

NBER WORKING PAPER SERIES

LEARNING, CONFIDENCE, AND BUSINESS CYCLES

Cosmin L. Ilut
Hikaru Saijo

Working Paper 22958
<http://www.nber.org/papers/w22958>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2016

We would like to thank George-Marios Angeletos, Yan Bai, Nick Bloom, Fabrice Collard, Jesus Fernandez-Villaverde, Tatsuro Senga, Martin Schneider, Mathieu Taschereau-Dumouchel, Carl Walsh, as well as to seminar and conference participants at the "Ambiguity and its Implications in Finance and Macroeconomics" Workshop, Bank of England, Canon Institute for Global Studies, EIEF, Northwestern, NBER Summer Institute, San Francisco Fed, Stanford, UC Santa Cruz, SED Annual Meeting, the SITE conference on Macroeconomics of Uncertainty and Volatility, and Wharton for helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Cosmin L. Ilut and Hikaru Saijo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Learning, Confidence, and Business Cycles
Cosmin L. Ilut and Hikaru Saijo
NBER Working Paper No. 22958
December 2016
JEL No. C11,D81,E32,E44

ABSTRACT

We build a tractable heterogeneous-firm business cycle model where firms face Knightian uncertainty about their profitability and learn it through production. The cross-sectional mean of firm-level uncertainty is high in recessions because firms invest and hire less. The higher uncertainty reduces agents' confidence and further discourages economic activity. We characterize this feedback mechanism in linear, workhorse macroeconomic models and find that it endogenously generates empirically desirable cross-equation restrictions such as: amplified and hump-shaped dynamics, co-movement driven by demand shocks and countercyclical correlated wedges in the equilibrium conditions for labor, risk-free and risky assets. In a rich model estimated on US macroeconomic and financial data, the information friction changes inference and significantly reduces the empirical need for standard real and nominal rigidities. Furthermore, endogenous idiosyncratic uncertainty propagates shocks to financial conditions, disciplined by observed spreads, as key drivers of fluctuations, and magnifies the aggregate activity's response to monetary and fiscal policies.

Cosmin L. Ilut
Department of Economics
Duke University
223 Social Sciences Building
Box 90097
Durham, NC 27708
and NBER
cosmin.ilut@duke.edu

Hikaru Saijo
University of California at Santa Cruz
Economics Department
401 Engineering 2 Building
1156 High Street
Santa Cruz, CA 95064
hsaijo@ucsc.edu

1 Introduction

A fundamental issue in the business cycle analysis is to model the source of frictions that matter for the economy's response to shocks. On the one hand, for a given set of structural shocks and policy actions, a specific propagation mechanism provides the identification needed to evaluate the relative importance of disturbances in explaining the observed macroeconomic data. On the other hand, a change in the nature of shocks or in the policy actions leads to different economic outcomes depending on the frictions that are taken to be policy-invariant.

In this paper we argue that imperfect observability about firm's profitability is an information friction that has the theoretical and quantitative potential to act as a key internal propagation mechanism in business cycle models that fit macroeconomic data well. In our model, there are two crucial aspects of uncertainty that matter for the mechanism. First, learning occurs through production. In particular, our environment implies that by producing at a larger scale, the firm reduces the variability of an unobserved temporary 'firm luck' component in its profitability and thus gets to observe a more precise signal about its unobserved persistent 'firm quality' component.¹ The second feature is that agents perceive uncertainty not just as risk but also as ambiguity, described here by the recursive multiple priors preferences.² This preference representation makes agents act *as if* they evaluate plans according to a worst case scenario drawn from a set of multiple beliefs. A wider set of beliefs corresponds to a loss of confidence in probability assessments. In particular, we assume that when facing a larger estimation uncertainty, the agent is less confident about the conditional mean of the underlying persistent firm quality component. The lower confidence makes the agent behave as if the worst-case mean becomes worse.³

The two ingredients generate a feedback loop at the firm level: lower production leads to more estimation uncertainty, which in turn shrinks the optimal size of productive inputs. In our model, the firm-level feedback loop linearly aggregates so that recessions are periods of a high cross-sectional mean of firm-level uncertainty because firms on average invest and hire less. In turn, the high average uncertainty further dampens aggregate activity. We show how this feedback turns our heterogeneous firm model into a standard linear business cycle model with a modified equilibrium law of motion compared to its rational expectations version.

We characterize the equilibrium dynamics with endogenous uncertainty and find strong internal

¹What matters for our mechanism is the procyclical signal to noise ratio at firm level. We discuss several modeling approaches to generate this, which produce the same dynamics at the aggregate. Our specific approach is to consider two unobservable shocks into a firm's production. Each unit of production is hit by a persistent firm-specific productivity shock, which we call 'firm quality', and by a transitory, unit-specific shock, uncorrelated with firm quality and independently distributed over units in the economy. The firm's total observed output is thus given by the observed number of production units times their observed average productivity. This average equals the sum of unobserved firm quality and of the unobserved average realization of the unit-specific shocks, which we call 'firm luck'. It follows that this luck component is more variable the lower is the number of units produced.

²The multiple priors preferences have been introduced by Gilboa and Schmeidler (1989). Epstein and Schneider (2003b) provide axiomatic foundations for the extension to intertemporal choice.

³This is simply a manifestation of precautionary behavior, which lowers the certainty equivalent of the return to production, but, compared to risk, it allows for first-order effects of uncertainty on decisions.

propagation, with amplified and hump-shaped responses to shocks. In addition, the equilibrium law of motion is characterized by the following cross-equation restrictions: (i) endogenous countercyclical wedges in the labor optimality condition as well as in the pricing of a risk-free bond and a risky asset such as capital; (ii) positive co-movement of aggregate variables in response to either supply or demand shocks. When the model is fit to a set of observed macroeconomic and financial data, these cross-equation restrictions provide over-identifying discipline on the mechanism.

Propagation mechanism. The feedback between the agent's degree of confidence and the endogenous state of the economy results in an implied 'as if' confidence process that sustains the equilibrium allocation. Fluctuations in confidence lead to state-dependent wedges between the worst-case belief and the econometrician's data generating process (DGP). For example, in a recession, caused either by supply or demand shocks, the lower degree of confidence widens the belief distortion. The distortion appears in state prices and therefore in all the relevant Euler equations. As such, the model produces correlated, countercyclical wedges. In turn, the state-dependent distortion affects dynamics through those Euler equations. For example, in a recession, the lower confidence decreases the uncertainty-adjusted return to labor, consumption and capital which reduces incentives to work, consume and invest. Of particular importance for macroeconomic aggregates is the lower perceived return to hiring, manifested as a high labor wedge, which leads to low equilibrium hours worked even if consumption is low and the realized marginal product of labor is on average unchanged under the econometrician's DGP.⁴ By producing an endogenous labor wedge, the model can generate co-movement between consumption, hours and investment as a response to either supply or demand shocks.

We analyze the mechanism of endogenous confidence by first exploring a basic RBC model that besides imperfect information is otherwise frictionless. This provides a useful laboratory to explain the formal and intuitive arguments on how the model is solved with linear methods as well as the mechanisms through which uncertainty maps into correlated endogenous wedges.

Importantly for our aggregation result, we show that when uncertainty is modeled as ambiguity with multiple priors that differ in their mean, the law of large numbers implies that sets of beliefs converge, but that idiosyncratic uncertainty cannot be diversified away. Such results have been developed in decision-theoretical work on the law of large numbers under ambiguity, as in Marinacci (1999) and Epstein and Schneider (2003a), but have not yet been used in macroeconomic models.

More specifically, we model an otherwise typical representative agent that works, consumes the final good, trades bonds and owns the portfolio of firms. This agent therefore has to forecast the aggregate dividend of this portfolio. In models in which uncertainty over the firms' dividends is modeled only as risk, the agent is assumed to have full confidence in the probability distributions of firms' productivity. This results in a law of large numbers that eliminates the role of idiosyncratic uncertainty. The only difference from the standard model is that our agent is not confident about the distributions of firms' individual productivity and hence profit. Faced with this ambiguity, the

⁴See Shimer (2009) and Chari et al. (2007) for evidence and discussion of labor wedges.

agent evaluates choices *as if* mean firm productivity is low. Therefore, the representative agent who owns a portfolio of many firms acts *as if* the entire portfolio has a low mean payoff. This means that the confidence about idiosyncratic conditions does not vanish in the aggregate.

We highlight the key parameters that control the feedback between uncertainty and economic activity. Since uncertainty is an equilibrium object in our model, in general the whole set of structural parameters matters for this feedback. It is useful to separate these in three categories. First, the standard production and preference parameters that control the variability of the production inputs, such as the Frisch elasticity or capital share. The second set includes the persistence of firm-level shocks and the steady state signal-to-noise ratio for the firm’s learning problem. For the former we use values from the empirical literature on micro-level data, such as Kehrig (2015). The latter parameter controls not just the average amount of estimation uncertainty, but also how variation in the productive inputs maps into the uncertainty variation. While quantifying the learning friction at the firm level is inherently difficult, we appeal to recent empirical work by David et al. (2015), who estimate a firm-level signal-to-noise ratio relevant for our model. The third key parameter is the entropy constraint, which determines not only the average ambiguity level, but also how fluctuations in the estimation uncertainty map into fluctuations in ambiguity. To discipline this parameter, we use a model consistency criterion developed in Ilut and Schneider (2014), according to which agents’ ambiguity should not be “too large”, in a statistical sense, compared to the variability of the data.

We compare the behavior of the information friction model against the version where this is absent. This comparison is controlled by only one parameter, the entropy constraint. Indeed, in the rational expectations model the effects of firm-level learning disappear entirely due to linearization and the law of large numbers. Our impulse response analysis highlights the key forces that the friction brings. First, we find that endogenous uncertainty is a powerful propagation mechanism. A positive shock that raises economic activity increases the level of confidence, which in turn further affects economic activity, leading to an amplified and hump-shaped impulse response.

Second, we explain how endogenous uncertainty generates countercyclical correlated wedges, a feature that is strongly present in the data. In our model, in periods of low production, when estimation uncertainty is larger, firms lower its demand for inputs ‘excessively’, as compared to what the econometrician measures for the equilibrium marginal product of labor or realized return on capital. Because uncertainty has first order effects on firms’ decision rules these reduced-form wedges linearly aggregate up. Thus, during recessions, we observe an unusually low level of equilibrium hours, manifested as a higher labor wedge. The endogenous uncertainty also leads to a countercyclical desire to save in risk-free assets. From the perspective of an econometrician that measures the realized growth rate of marginal utility, the low real rates in recessions look like countercyclical discount factor wedges.⁵ Finally, the increase in ambiguity also makes capital

⁵See Christiano et al. (2005) and Smets and Wouters (2007) as examples for the large literature on estimated DSGE models with shocks to the discount factor. In addition, the literature on the zero lower bound on nominal interest rate, such as Eggertsson et al. (2003) and Christiano et al. (2015), argue that disturbances to the demand for risk-free assets are key stochastic forces.

less attractive to hold. Investors holding an ambiguous asset are thus compensated by the higher measured excess return.⁶ The countercyclical labor, discount factor and financial wedges do not arise from separate shocks to labor supply, demand for risk-free assets and premia, or even from exogenous confidence shocks, but instead from any underlying shock that moves the economy.

Third, the model generates positive co-movement in response to shocks that do not move actual labor productivity or labor supply. Consider, for example, a negative shock to the expected return to investment, a form of what has been recently proposed in the literature as a reduced form of a negative financial shock. This shock produces standard intertemporal and wealth effects on consumption, which overall tends to lead to a fall in investment and hours, as well as a rise in consumption. In standard RBC models, consumption and hours co-move positively only in the long run after the negative productivity effect of capital accumulation kicks in. In our model, the lower economic activity decreases average confidence. This leads to a stronger negative wealth effect, as well as a lower perceived marginal return to labor. The negative effect can be strong enough to make consumption fall and the labor wedge can be high enough to simultaneously lead to low equilibrium labor. Therefore, our model can generate positive co-movement in consumption and labor after a demand shock much quicker than the typical effects present in standard models that arise from endogenous capital accumulation.

Fourth, firms have an incentive to actively use their productive inputs to learn about their profitability. We show that this experimentation incentive has a first-order effect in our setup and we characterize its cyclical properties. We find that experimentation exerts a procyclical influence on aggregates in the short run, turning experimentation into an amplifying factor, and a countercyclical one in the medium term, where it dampens fluctuations.

Quantitative results. We use the linearity of the solution method to estimate by likelihood based methods a quantitative model with a large state space. Our estimated model includes an aggregate TFP shock and a financial shock, analyzed for example in Christiano et al. (2015). We use standard observables for US aggregate data: the growth rate of output, hours worked, investment and consumption, as well as inflation and nominal interest rate. In addition, we use the Baa corporate spread as an observable proxy for the financial shock.

The estimated model provides evidence that endogenous confidence significantly changes the inference on the role of rigidities and shocks driving fluctuations. In terms of rigidities, learning reduces the need of additional frictions for fitting the data. In particular, compared to the rational expectations (RE) version, the habit formation parameter is lowered by 40%, the investment adjustment cost becomes negligible, the average Calvo adjustment period of prices falls from 9 to 2 quarters, and for wages from 6 quarters to 1 quarter. The reason for these smaller estimated frictions is that the learning mechanism provides strong internal propagation and induces by itself positive co-movement in response to the financial shock.

In terms of model comparison, the model with endogenous uncertainty fits the data significantly

⁶For a review of the evidence on countercyclical excess returns see Cochrane (2011).

better. The reason is twofold. First, as described above, learning emerges as a parsimonious friction, as opposed to the array of other rigidities. Second, in the RE version the observed spread generates a counterfactual negative co-movement between observed macroeconomic variables. The information friction changes significantly the propagation mechanism to the point that the observed countercyclical spread is now consistent with co-movement and thus the estimation prefers observed contractionary financial shocks during recessions. At the same time, in our model agents lose confidence during recessions, which further contributes to the countercyclicality of the spread endogenously and therefore improves the model fit.

Our quantitative results point to an important takeaway on the role of financial shocks. On the one hand, part of the observed spread is explained by the endogenous model-implied component and therefore there is less empirical need for finding financial shocks. On the other hand, the propagation mechanism produced by our information friction makes the estimated shocks, whose time-series properties are disciplined by the observable proxy, have large and empirically plausible effects on the observed macroeconomic aggregates.

Finally, while the feedback between activity and uncertainty puts strong discipline on the estimated confidence process, we further test the model implied path for ambiguity against an observable that has not been used for the estimation. In particular, in our model the lack of confidence about the probability distribution of firm's profitability is reflected in a set of conditional mean forecasts about each firm's return on capital. We use a recent time-series dataset on survey forecasts by financial analysts built by Senga (2015) that refers to the cross-sectional average dispersion of conditional mean forecasts about firm-level returns. We find that the time-series path for the model-implied dispersion matches the empirical counterpart very well.

Relation to literature. We relate to a recent literature on confidence shocks. Indeed, once our equilibrium 'as if' confidence process is taken as given, the economic mechanisms by which confidence impacts decisions through Euler equation wedges are therefore common to models with exogenous confidence shocks. In this context, Ilut and Schneider (2014) allow for time-variation in confidence about aggregate conditions that arises from exogenous ambiguity, while Angeletos and La'O (2009, 2013) and Angeletos et al. (2014) study confidence shocks in the form of correlated higher order beliefs in models of dispersed information. This work has analyzed the impact of belief shocks on Euler equations and the quantitative appeal of such confidence shocks in fitting macroeconomic dynamics. In particular, Angeletos et al. (2014) describe desirable properties of these beliefs shocks, such as the focus on short-run rather than long-run activity, and their empirical success. Through these lenses, the 'as if' confidence interpretation of our model helps explain why the information friction model fits the data well. Indeed, the friction provides the endogenous mechanism to map the observed financial shocks into countercyclical movements in confidence about short-run activity, which in turn, at their equilibrium path, affects macroeconomic dynamics through empirically desirable countercyclical wedges.

By connecting to the existing literature on exogenous confidence shocks, our approach is

complementary, as it provides a propagation mechanism for how these processes can arise. As a consequence, we put significant discipline on the nature and size of the confidence shocks. In particular, the parameters that matter for confidence have specific interpretation since we provide a micro-founded model. In addition, the information friction ties, through the endogenous equilibrium law of motion, the 'as if' confidence process to the observed macroeconomic and financial data. For policy purposes, the micro-foundation also makes clear that the 'as if' confidence process cannot be taken as policy-invariant since it depends on the whole equilibrium path of actions.

Our paper is also related to a rapidly growing literature on time-varying uncertainty. There are models that generate recessions through exogenous increases in the volatility of firm-level or aggregate productivity.⁷ Compared to that literature, we do not rely on additional real or nominal rigidities to generate contractionary effects of higher uncertainty and positive co-movement. The reason is that in our model the level of input (such as labor) is chosen under imperfect information, and thus under uncertainty, about the underlying productivity process. This leads in our model to a countercyclical labor wedge, an issue that we take on below. In addition, in line with the 'as if' confidence argument above, in our economy changes in perceived uncertainty are endogenous and hence affect the propagation of any aggregate shocks, including exogenous volatility or confidence shocks, as well as outcomes of policy experiments.

An important outcome of our proposed mechanism is the endogenous labor wedge. It is useful to observe that in a traditional model labor supply is chosen based on the same information set as consumption and the return to working is perfectly observed when the labor demand is made. For belief distortions or time-varying uncertainty to matter, the literature has deviated from that environment by typically allowing for a labor choice made under uncertainty about its return. Angeletos and La'O (2009, 2013) describe the key income and substitution forces through which confidence shocks show up as labor wedges in a model where hiring occurs under such uncertainty. Angeletos et al. (2014) further emphasize the critical role of beliefs being about the short-run rather than the long-run activity in producing stronger substitution effects.

In the literature, this 'labor in advance' timing is typically assumed as part of a technological friction,⁸ or it arises as a manifestation of an information friction. The latter category includes models where the individual firm needs to forecast the actions of others, such as in models of dispersed information referenced above, or in models where firms learn from an endogenously determined action, such as production in our model and a learning literature described below. Indeed, in our context, the information set used to learn about exogenous fundamentals conditions

⁷For idiosyncratic volatility see for example Bloom (2009), Arellano et al. (2012), Bachmann and Bayer (2013) and Christiano et al. (2014). For aggregate volatility see for example Fernández-Villaverde et al. (2011); Fernández-Villaverde et al. (2015), Basu and Bundick (2012), Born and Pfeifer (2014) and Bianchi et al. (2014).

⁸This timing assumption is typically made in models with some financial friction, where labor is assumed to be chosen before a cash flow shock is realized. In those models, a higher uncertainty, either exogenous (as in Arellano et al. (2012)) or endogenous (as in Gourio (2014), Straub and Ulbricht (2016)), about that cash flow realization, may lead to a labor wedge. Kozlowski et al. (2015) obtain a labor wedge instead from persistent belief distortions that arise from non-parametric learning about the distribution of shocks.

on optimally chosen actions, including labor. We therefore naturally obtain a lack of synchronization of labor and spending choices, a feature that Angeletos et al. (2014) discuss as key for allowing confidence movements to manifest as labor wedges.

The low scale - high uncertainty feedback characteristic of our model is present in different forms in a literature on learning. The typical approaches, as in Caplin and Leahy (1993), van Nieuwerburgh and Veldkamp (2006), Ordoñez (2013) and Fajgelbaum et al. (2016), have two main common features: (i) the feedback matters through non-linear dynamics and (ii) uncertainty is about aggregate conditions.⁹ Here we show that expanding the notion of uncertainty to ambiguity allows studying an endogenous uncertainty mechanism in models with: (i) tractable linear dynamics and (ii) uncertainty about idiosyncratic conditions.

Our approach advances the modeling of information frictions in business cycle models in two major ways, which connect back to our fundamental motivation of building micro-founded structural models that are useful for policy interventions. The first is quantitative: why and by how much does the friction matter empirically? Working with tractable linear dynamics allows the information friction to be embedded in quantitative macroeconomic models that use standard solution and estimation methods. This leads to a transparent characterization of the mechanism through which endogenous uncertainty matters for key business cycle patterns, such as the delayed and persistent response to shocks and the 'as if' countercyclical wedges.

Importantly, we are able to evaluate its quantitative effect by empirically fitting a standard set of macroeconomic variables even within models that include widely-used alternative frictions, such as nominal or real rigidities. To generate persistent and hump-shaped dynamics, arbitrary frictions such as consumption habit and investment adjustment costs are often added to the standard models. Sticky prices and wages are also usually important because they generate additional persistence and help break the Barro and King (1984) critique to make other types of shocks, besides productivity or intratemporal labor supply shocks, generate positive co-movements of macro aggregates.¹⁰ In our estimated model, we find that endogenous uncertainty generates a rich set of empirically desirable cross-equation restrictions that reduce significantly the need for such additional rigidities. At the same time, these restrictions affect the identification of the relevant disturbances. In particular, endogenous uncertainty changes the propagation mechanism to the point that financial shocks, disciplined by observed asset prices in our estimated model, can generate a positive co-movement in macroeconomic aggregates and thus become a significant driver of business cycle dynamics.

⁹The reasons why uncertainty matters in that literature differ, including the representative agent's risk aversion, financial frictions or irreversible investment costs. As with confidence shocks in linear models, a countercyclical labor wedge arises in these non-linear models if the substitution effect from a lower risk-adjusted return to working overcomes the income effect. Evaluating those wedges has not been typically the object of those models.

¹⁰Recent line of attacks to break Barro and King (1984) critique include: strategic complementarity in a model with dispersed information (Angeletos and La'O (2013), Angeletos et al. (2014)), heterogeneity in labor supply and consumption across employed and non-employed (Eusepi and Preston (2015)), variable capacity utilization together with a large preference complementarity between consumption and hours (Jaimovich and Rebelo (2009)), and the large literature on countercyclical markups through nominal rigidities.

The second, and related, question is that of policy interventions: how does policy affect the economy differently in the information friction model? On the one hand, the source of uncertainty, whether it is about firm-level or aggregate conditions, matters for policy recommendations. In our model, linearity allows for an easy aggregation of decision rules and therefore for a tractable modeling of uncertainty about firm-level shocks. This not only extends the plausible sources of imperfect information to firm-level volatilities, which are empirically larger than aggregate fluctuations, but also means that the competitive equilibrium is constrained Pareto optimal.¹¹ Thus, differently from learning about aggregate conditions, in this world there are no information externalities since learning occurs at the individual firm level and not from observing the aggregate economy.¹²

On the other hand, the endogeneity of uncertainty matters because it transmits policy changes differently compared to an exogeneity benchmark. We show that in our estimated model a policy change where the Taylor rule responds to the financial spread would significantly lower output variability only because it stabilizes the endogenous variation in uncertainty.¹³ For fiscal policy we find a significantly larger government spending multiplier because of the effect on confidence.

The paper is structured as follows. In Section 2 we introduce our heterogeneous-firm model and discuss the solution method. We describe the endogenous countercyclical wedges and illustrate its dynamic properties through impulse response analysis in Section 3. In Section 4 we add additional rigidities to estimate a model on US aggregate data.

2 The model

Our baseline model is a real business cycle model in which, as in the standard framework, firms are owned by the household and maximize shareholder value. We augment the standard framework with two key features: the infinitely-lived representative household is ambiguity averse and that ambiguity is about the firm-level profitability processes.

2.1 Technology

There is a continuum of firms, indexed by $l \in [0, 1]$, which act in a perfectly competitive manner. They use capital $K_{l,t-1}$ and hire labor $H_{l,t}$ to produce $\tilde{Y}_{l,t}$ units of good l by a production function

$$\tilde{Y}_{l,t} = K_{l,t-1}^\alpha H_{l,t}^{1-\alpha} \quad (2.1)$$

A firm that produces good l consists of production units, indexed by j . These units are defined

¹¹Our approach can incorporate learning not only about idiosyncratic but also aggregate conditions, and as such, our results can be viewed as a lower bound on the overall business cycle effects of endogenous uncertainty.

¹²For example, the increased economic activity, and the associated average reduction in uncertainty, produced by a government spending increase can not be welfare increasing in our model.

¹³Such a policy was proposed by Taylor (2008), among others. Cúrdia and Woodford (2010) and Christiano et al. (2010) study the effects of similar policies using standard DSGE models.

such that combining capital $K_{j,l,t-1}$ and labor $H_{j,l,t}$ results in one unit of good l . The firm obtains more production by operating more such units. Each unit produces a stochastic quality adjusted output, which is driven by three components: an aggregate shock, a firm-specific shock and an idiosyncratic shock. In particular, this output equals

$$A_t (z_{l,t} + \tilde{\nu}_{l,j,t}), \quad (2.2)$$

where A_t is an aggregate technology shock that follows

$$\ln A_t = \rho_A \ln A_{t-1} + \epsilon_{A,t}, \quad \epsilon_{A,t} \sim N(0, \sigma_A^2), \quad (2.3)$$

$z_{l,t}$ is a firm-specific technology shock that follows¹⁴

$$z_{l,t} = (1 - \rho_z)\bar{z} + \rho_z z_{l,t-1} + \epsilon_{z,l,t}, \quad \epsilon_{z,l,t} \sim N(0, \sigma_z^2), \quad (2.4)$$

and the unit specific shock is iid $\tilde{\nu}_{l,j,t} \sim N(0, \sigma_\nu^2)$.

Since the firm is able to operate $\tilde{Y}_{l,t}$ number of units given by (2.1) and each unit produces according to (2.2), the firm's total output equals

$$Y_{l,t} = \tilde{Y}_{l,t} A_t (z_{l,t} + \nu_{l,t}), \quad (2.5)$$

where the average realization over unit specific shocks, $\nu_{l,t}$, is given by

$$\nu_{l,t} \equiv \frac{1}{\tilde{Y}_{l,t}} \sum_{j=1}^{\tilde{Y}_{l,t}} \tilde{\nu}_{l,j,t} \sim N\left(0, \frac{\tilde{Y}_{l,t} \sigma_\nu^2}{\tilde{Y}_{l,t}^2}\right) \quad (2.6)$$

Given production outcomes and its associated costs, firms pay out dividends

$$D_{l,t} = Y_{l,t} - W_t H_{l,t} - I_{l,t}, \quad (2.7)$$

where W_t is the real wage and $I_{l,t}$ is investment. Their capital stock follows the law of motion

$$K_{l,t} = (1 - \delta)K_{l,t-1} + I_{l,t}.$$

2.2 Imperfect information

Agents cannot directly observe the realizations of firm-specific shocks $z_{l,t}$ and $\tilde{\nu}_{l,j,t}$. Instead, every agent in the economy observes the aggregate shocks, the inputs used for the production of units $\tilde{Y}_{l,t}$, as well as the quality adjusted output $Y_{l,t}$ of each firm l . As a consequence, from the production

¹⁴We assume that $z_{l,t}$ follows a normal process, a technical assumption that is useful in solving the learning problem because it makes the information friction linear.

function in (2.5), agents observe the sum $(z_{l,t} + \nu_{l,t})$, but not its individual components. This informational assumption leads to a non-invertibility problem: Agents cannot tell whether an unexpectedly high realization of firm's output $Y_{l,t}$ is due to the firm being 'better' (an increase in the persistent firm's specific profitability $z_{l,t}$) or just 'lucky' (an increase in the average realization of the unit specific shocks $\nu_{l,t}$).¹⁵

Faced with this uncertainty, agents use the available information, including the path of output and inputs, to form estimates on the underlying source of profitability $z_{l,t}$. Since the problem is linear and Gaussian, Bayesian updating using Kalman filter is optimal from the statistical perspective of minimizing the mean square error of the estimates.¹⁶ After production at period t , the measurement equation of the Kalman filter is given by

$$Y_{l,t}/(\tilde{Y}_{l,t}A_t) = z_{l,t} + \nu_{l,t}, \quad \nu_{l,t} \sim N\left(0, \tilde{Y}_{l,t}^{-1}\sigma_\nu^2\right)$$

and the transition equation is given by equation (2.4).

Note that, unlike the standard time-invariant Kalman filter, the variance in the measurement equation, $\tilde{Y}_{l,t}^{-1}\sigma_\nu^2$, is time-varying.¹⁷ This reflects that when the firm l produces more units $\tilde{Y}_{l,t}$ the 'luck' component, $\nu_{l,t}$, becomes less variable. Indeed, as shown by equation (2.6), the firm knows that the average of the unit specific shocks is a random variable with a variance that decreases in the number of units produced. As such, when $\tilde{Y}_{l,t}$ is higher the firm extracts the information that the observed firm output $Y_{l,t}$ is more likely to have been generated by the firm specific profitability shock $z_{l,t}$. Put differently, the signal-to-noise ratio becomes procyclical.¹⁸

The flip side implication of the procyclical signal-to-noise ratio is that uncertainty is countercyclical: the posterior variance of idiosyncratic technology $z_{l,t}$ rises during recessions. Intuitively, when a firm puts less resources into production, its estimate about its productivity is imprecise because the level of output is largely determined by the realization of the transitory shock. Conversely, its estimate is accurate when it uses more resources because output mostly reflects the realization of productivity $z_{l,t}$.

To characterize the filtering problem, we start by deriving the one-step-ahead prediction from

¹⁵The revenue obtained by firm l is $P_t Y_{l,t}$, where we normalize $P_t = 1$. We can rewrite the revenue as $P_{l,t} \tilde{Y}_{l,t}$, where the relative price paid for the units of good l , $P_{l,t}$, is given by its quality adjustment $Y_{l,t}/\tilde{Y}_{l,t}$. This shows that the shocks $z_{l,t}$ and $\tilde{\nu}_{l,j,t}$ can be equivalently interpreted as productivity or quality shocks, and so we use them interchangeably as profitability shocks.

¹⁶The learning problem of the model with growth is in Appendix 6.3, along with other equilibrium conditions.

¹⁷Note that, after production, this variance is pre-determined, which allows us to use the Kalman filter.

¹⁸The pro-cyclicality of the signal-to-noise ratio is obtained as long as the unit specific shock $\tilde{\nu}_{l,j,t}$ is not perfectly correlated with the firm specific shock. With less than perfect correlation, there is still residual uncertainty about the realization of the average $\nu_{l,t}$ that decreases with the number of units $\tilde{Y}_{l,t}$.

the period $t - 1$ estimate $\tilde{z}_{l,t-1|t-1}$ and its associated error variance $\Sigma_{l,t-1|t-1}$. We have

$$\begin{aligned}\tilde{z}_{l,t|t-1} &= (1 - \rho_z)\bar{z} + \rho_z\tilde{z}_{l,t-1|t-1}, \\ \Sigma_{l,t|t-1} &= \rho_z^2\Sigma_{l,t-1|t-1} + \sigma_z^2.\end{aligned}$$

Then, given observables $(Y_{l,t}, \tilde{Y}_{l,t}$ and $A_t)$ firms update their estimates according to

$$\tilde{z}_{l,t|t} = \tilde{z}_{l,t|t-1} + \text{Gain}_{l,t}(Y_{l,t}/(\tilde{Y}_{l,t}A_t) - \tilde{z}_{l,t|t-1}), \quad (2.8)$$

where $\text{Gain}_{l,t}$ is the Kalman gain and is given by

$$\text{Gain}_{l,t} = \left[\frac{\Sigma_{l,t|t-1}}{\Sigma_{l,t|t-1} + \tilde{Y}_{l,t}^{-1}\sigma_\nu^2} \right]. \quad (2.9)$$

The updating rule for variance is

$$\Sigma_{l,t|t} = \left[\frac{\sigma_\nu^2}{\tilde{Y}_{l,t}\Sigma_{l,t|t-1} + \sigma_\nu^2} \right] \Sigma_{l,t|t-1}. \quad (2.10)$$

Intuitively, the error variance is increasing in the un-informativeness of the observation, which is the variance of noise divided by the total variance. We can see that, holding $\Sigma_{l,t|t-1}$ constant, the posterior variance $\Sigma_{l,t|t}$ increases as input $\tilde{Y}_{l,t}$ decreases.

The dynamics of the idiosyncratic technology $z_{l,t}$ according to the Kalman filter can thus be described as

$$z_{l,t+1} = (1 - \rho_z)\bar{z} + \rho_z(\tilde{z}_{l,t|t} + u_{l,t}) + \epsilon_{z,l,t+1}, \quad (2.11)$$

where $u_{l,t}$ is the estimation error of $z_{l,t}$ and $u_{l,t} \sim N(0, \Sigma_{l,t|t})$.

2.3 Household wealth

There is a representative household whose budget constraint is given by

$$C_t + B_t + \int P_{l,t}^e \theta_{l,t} dl \leq W_t H_t + R_{t-1} B_{t-1} + \int (D_{l,t} + P_{l,t}^e) \theta_{l,t-1} dl + T_t,$$

where B_t is the one-period riskless bond, W_t is the real wage, R_t is the interest rate, and T_t is a transfer. $D_{l,t}$ and $P_{l,t}^e$ are the dividend and price of a unit of share $\theta_{l,t}$ of firm l , respectively.

Market clearing and resource constraint

We impose the market clearing conditions for the labor market and the bond market:

$$H_t = \int_0^1 H_{l,t} dl, \quad B_t = 0.$$

The resource constraint is given by

$$C_t + I_t + G_t = Y_t \quad (2.12)$$

where $I_t \equiv \int_0^1 I_{l,t} dl$, G_t is the government spending and we assume a balanced budget each period ($G_t = -T_t$). We assume the government spending as a constant share of output: $\bar{g} = G_t/Y_t$.

A financial shock

On purpose, our model has no other frictions except the incomplete information about the firms' profitability. In particular, on the financing side, the model has a frictionless financial market. However, guided by recent theoretical and empirical developments, we are interested in analyzing the role of shocks to financial conditions as potential drivers of business cycle dynamics. In the context of our frictionless model, we follow existing reduced-form approaches in modeling such disturbances as a "financial wedge" shock Δ_t^k to agents' Euler equation for capital accumulation. In particular, as in Christiano et al. (2015), we introduce a shock Δ_t^k to the expected return on capital R_{t+1}^k such that

$$1 = (1 - \Delta_t^k) E_t^* M_{t+1} R_{t+1}^k.$$

Here $M_{t+1} \equiv M_0^{t+1}/M_0^t$, where M_0^t denote prices of t -period ahead contingent claims given by

$$M_0^t = \beta^t \lambda_t, \quad (2.13)$$

where λ_t is the marginal utility of consumption at time t by the representative household. Importantly for the propagation mechanism through the learning friction, these prices are evaluated under the representative agent's equilibrium conditional probability distribution E_t^* . We describe the properties of these beliefs in section 2.4 below.

We assume the process for Δ_t^k :

$$\Delta_t^k = (1 - \rho_\Delta) \bar{\Delta}^k + \rho_\Delta \Delta_{t-1}^k + \epsilon_{\Delta,t}, \quad \epsilon_{\Delta,t} \sim N(0, \sigma_\Delta^2). \quad (2.14)$$

As discussed in Christiano et al. (2015), the financial wedge could reflect variations in costs of financial intermediation such as bankruptcy costs or changes in the desirability of corporate bonds due to, for example, liquidity concerns.¹⁹ In our full-fledged estimated business cycle model of section 4, we will use observed variation in corporate bond spreads to discipline movements in this shock. In addition, as we will show later in section 3.3, our model generates additional movements in that spread through the endogenous variation in confidence. For now, in the baseline frictionless model developed here, our focus is to explore how endogenous uncertainty affects the propagation of the financial shock.

¹⁹The only distinction between the two interpretations is in their effect on the resource constraint. As in Christiano et al. (2015), the quantitative magnitude of this impact is likely to be very small and hence we have ignored it for simplicity in our market clearing condition (2.12).

2.4 Optimization

We have described so far the firms' production possibilities, the household budget constraint, the sources of stochasticity and the available information set. We now present the optimization problems of the representative household and of the firms.

Imperfect information and ambiguity

The representative household perceives ambiguity (Knightian uncertainty) about the vector of idiosyncratic productivities $\{z_{l,t}\}_{l \in [0,1]}$. We now describe how that ambiguity process evolves.

The agent uses observed data to learn about the hidden technology by using the Kalman filter to obtain a benchmark probability distribution. The Kalman filter problem has been described in section 2.2. Ambiguity is modeled as a one-step ahead set of conditional beliefs that consists of alternative probability distributions surrounding the benchmark Kalman filter estimate $\tilde{z}_{l,t}$ in (2.11) of the form

$$z_{l,t+1} = (1 - \rho_z)\bar{z} + \rho_z\tilde{z}_{l,t|t} + \mu_{l,t} + \rho_z u_{l,t} + \epsilon_{z,l,t+1}, \quad \mu_{l,t} \in [-a_{l,t}, a_{l,t}] \quad (2.15)$$

In particular, the agent considers a set of alternative probability distributions surrounding the benchmark that is controlled by a bound on the relative entropy distance. More precisely, the agent only considers the conditional means $\mu_{l,t}$ that are sufficiently close to the long run average of zero in the sense of relative entropy:

$$\frac{\mu_{l,t}^2}{2\rho_z^2\Sigma_{l,t|t}} \leq \frac{1}{2}\eta^2, \quad (2.16)$$

where the left hand side is the relative entropy between two normal distributions that share the same variance $\rho_z^2\Sigma_{l,t|t}$, but have different means ($\mu_{l,t}$ and zero), and η is a parameter that controls the size of the entropy constraint. The entropy constraint (2.16) results in a set $[-a_{l,t}, a_{l,t}]$ for $\mu_{l,t}$ in (2.15) that is given by

$$a_{l,t} = \eta\rho_z\sqrt{\Sigma_{l,t|t}} \quad (2.17)$$

The interpretation of the entropy constraint is that agents are less confident, i.e. the set of beliefs is larger, when there is more estimation uncertainty. The relative entropy can be thought of as a measure of distance between the two distributions. When uncertainty $\Sigma_{l,t|t}$ is high, it becomes difficult to distinguish between different processes. As a result, agents become less confident and contemplate wider sets $\mu_{l,t}$ of conditional probabilities. Therefore, from the perspective of the agent, a change in posterior variance translates into a change in uncertainty about the one-step-ahead mean realization of technology.

Household problem

We model the household's aversion to ambiguity through recursive multiple priors preferences, which capture an agent's lack of confidence in probability assessments. This lack of confidence is manifested in the *set* of one step ahead conditional beliefs about each $z_{l,t+1}$ given in equations

(2.15) and (2.17). Collect the exogenous state variables in a vector $s_t \in S$. This vector includes the aggregate shocks A_t and Δ_t^k , as well as the cross-sectional distribution of idiosyncratic productivities $\{z_{l,t}\}_{l \in [0,1]}$. A household consumption plan C gives, for every history s^t , the consumption of the final good $C_t(s^t)$ and the amount of hours worked $H_t(s^t)$. For a given consumption plan C , the household recursive multiple priors utility is defined by

$$U_t(C; s^t) = \ln C_t - \frac{H_t^{1+\phi}}{1+\phi} + \beta \min_{\mu_{l,t} \in [-a_{l,t}, a_{l,t}], \forall l} E^\mu[U_{t+1}(C; s^t, s_{t+1})], \quad (2.18)$$

where β is the subjective discount factor and ϕ is the inverse of Frisch labor supply elasticity.²⁰

Notice that there is a cross-sectional distribution of sets of beliefs over the future $\{z_{l,t+1}\}_{l \in [0,1]}$. Indeed, for each firm l , the agent entertains a set of conditional means $\mu_{l,t} \in [-a_{l,t}, a_{l,t}]$. If each set is singleton we obtain standard separable log utility with those conditional beliefs. Otherwise, agents are not willing to integrate over the beliefs and narrow down each set to a singleton. In response, households take a cautious approach to decision making and act *as if* the true data generating process is given by the worst-case conditional belief.

Worst-case belief

The worst-case belief can be easily solved for at the equilibrium consumption plan: the worst-case expected idiosyncratic productivity is low. Given the bound in equation (2.17), the worst-case mean is

$$\mu_{l,t}^* = -\eta\rho_z \sqrt{\Sigma_{l,t|t}}. \quad (2.19)$$

Thus, the agent's cautious behavior faced with the sets of beliefs $\mu_{l,t} \in [-a_{l,t}, a_{l,t}]$, one for each l , results in acting *as if* the conditional mean of each firm's idiosyncratic productivity is given by the worst-case mean in (2.19). We denote conditional moments under these worst case belief by stars.

The worst-case conditional mean for each firm's $z_{l,t+1}$ is therefore given by

$$E_t^* z_{l,t+1} = (1 - \rho_z)\bar{z} + \rho_z \tilde{z}_{l,t|t} - \eta\rho_z \sqrt{\Sigma_{l,t|t}} \quad (2.20)$$

where $\tilde{z}_{l,t|t}$ is the Kalman filter estimate of the mean obtained in equation (2.8).

Law of large numbers

The worst-case conditional distribution of each firm's productivity is

$$z_{l,t+1} \sim N(E_t^* z_{l,t+1}, \rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) \quad (2.21)$$

Once the worst-case distribution is determined, it is easy to compute the cross-sectional average

²⁰The recursive formulation ensures that preferences are dynamically consistent. Details and axiomatic foundations are in Epstein and Schneider (2003b). Subjective expected utility obtains when the set of beliefs collapses to a singleton.

realization $\int z_{l,t+1} dl$. Indeed, by the law of large numbers this average converges to

$$\int E_t^* z_{l,t+1} dl = (1 - \rho_z) \bar{z} + \rho_z \int \tilde{z}_{l,t|t} dl - \eta \rho_z \int \sqrt{\Sigma_{l,t|t}} dl. \quad (2.22)$$

Moreover, it can be easily showed that $\int \tilde{z}_{l,t|t} dl = 0$. Intuitively, since under the true DGP the cross-sectional mean of $z_{l,t}$ is constant, the cross-sectional mean of the Kalman posterior mean estimate is a constant as well.

Equation (2.22) is a manifestation of the law of large numbers where uncertainty is modeled as ambiguity (Knightian uncertainty) with multiple priors that differ in their mean. In this setting, law of large numbers imply that sets of beliefs converge, but that idiosyncratic Knightian uncertainty cannot be diversified away. Intuitively, agents care about uncertainty because it affects their future continuation utility, as shown in equation (2.18). The representative agent derives wealth through the average dividend from the portfolio of firms, and the continuation utility is increasing in wealth. Therefore, when this representative agent does not know the distribution of firms' productivity, and hence their dividends, she evaluates choices *as if* mean productivity is low. An agent who owns a portfolio of many firms then acts *as if* the entire portfolio has a low mean payoff.

Modeling idiosyncratic uncertainty as both risk and ambiguity matters crucially for its effect on the law of large numbers. When modeled as risk, uncertainty is diversified away, following the standard argument that once we control for the conditional mean, given there by $\rho_z \tilde{z}_{l,t|t}$, the idiosyncratic variance $\Sigma_{l,t|t}$ does not matter for the cross-sectional average of expected $z_{l,t+1}$ in (2.22). Instead, with ambiguity, the variance $\Sigma_{l,t|t}$ also affects the set of mean beliefs. In particular, by (2.19), a larger estimation uncertainty makes it harder to distinguish between different processes, increasing ambiguity and therefore lowering the worst-case mean $E_t^* z_{l,t+1}$. Thus, the *average* idiosyncratic uncertainty, $\int \sqrt{\Sigma_{l,t|t}} dl$, matters for the *average* worst-case expected $z_{l,t+1}$ in (2.22).²¹

Firms' problem

Firms choose $K_{l,t}, H_{l,t}, I_{l,t}$ to maximize shareholder value

$$E_0^* \sum_{t=0}^{\infty} M_0^t D_{l,t}, \quad (2.23)$$

where E_0^* denotes expectation under the representative agent's worst case probability and $D_{l,t}$ is given by equation (2.7). The random variables M_0^t denote prices of t -period ahead contingent claims based on conditional worst case probabilities and are given by equation (2.13).

Compared to a standard model of full information and expected utility, the firm problem in (2.23) has two important specific characteristics. The first is that, as described above, unlike the

²¹See Marinacci (1999) or Epstein and Schneider (2003a) for formal treatments of the law of large numbers for i.i.d. ambiguous random variables. There they show that cross-sectional averages, such as that of equation (2.22), must (almost surely) lie in an interval bounded by the highest and lowest possible cross-sectional mean, and these bounds are tight in the sense that convergence to a narrower interval does not occur.

case of expected utility, the idiosyncratic uncertainty that shows up in these state prices does not vanish under diversification. The second concerns the role of experimentation. On one hand, with full information there is no role for experimentation. On the other hand, under incomplete information but Bayesian decision making, experimentation is valuable because it raises expected utility by improving posterior precision. Here, ambiguity-averse agents also value experimentation since it affects utility by tightening the set of conditional probability considered. Therefore, firms take into account in their problem (2.23) the impact of the level of input on worst-case mean. Although we allow such active learning by firms, our model can nevertheless still be solved using standard linear methods.²²

Role of information structure

The key property emerging from our information structure is the procyclical signal-to-noise ratio at the firm level. There are several other structural ways to generate this, which will produce the same dynamics at the aggregate. First, we could have an additive shock to the production function. In that case the produced output would simply be $Y_{l,t} = A_t(z_{l,t}\tilde{Y}_{l,t} + \gamma^t\nu_{l,t})$, where $\nu_{l,t}$ is an iid additive shock, unobserved at the firm level, with a constant variance. This setup is equivalent because it produces the same learning dynamics, in which a larger $\tilde{Y}_{l,t}$ means a larger coefficient in the measurement equation, and thus a higher signal-to-noise ratio.

Second, in Appendix 6.1, we offer an additional interpretation of this shock based on a more complex structure of imperfect competition. There we can refer to noisy demand signals and have a setup where firms learn more about the demand of their goods when they produce and sell more. In that version of the model, firms are subject to unobservable idiosyncratic shocks to the weight attached to their goods in the CES aggregator for final goods. There the temporary shock $\nu_{l,t}$ is replaced with an i.i.d. observation error of the underlying idiosyncratic shock; agents observe noisy signals about demand, whose precision is increasing in the level of individual production. Given the equivalent aggregate dynamics, we prefer the transparency of the current setup.

Timing

We find it useful to summarize the timing of events within a period t , which follows:

1. Stage 1 : Pre-production stage
 - Agents observe the realization of aggregate shocks (A_t and Δ_t^k).
 - Given forecasts about the idiosyncratic technology and its associated worst-case scenario, firms hire labor ($H_{l,t}$). The household supplies labor H_t and the labor market clears at the wage rate W_t .

2. Stage 2 : Post-production stage

²²When we present our quantitative results, we assess the contribution of experimentation by comparing our baseline results with those under passive learning, i.e. where agents do not actively experiment.

- Idiosyncratic shocks $z_{l,t}$ and $\nu_{l,t}$ realize (but are unobservable) and production takes place.
- Given output and input, firms update estimates about their idiosyncratic technology and use it to form forecasts for production next period.
- Firms make investment $I_{l,t}$ and pay out dividends $D_{l,t}$. The household makes consumption and asset purchase decisions (C_t , B_t , and $\theta_{l,t}$).

2.5 Log-linearized solution

We solve for the equilibrium law of motion using standard log-linear methods. This is possible for two reasons. First, since the filtering problem firms face is linear, the law of motion of the posterior variance can be characterized analytically (Saijo (2014)). Because the level of inputs has first-order effects on the level of posterior variance, linearization captures the impact of economic activity on confidence. Second, we consider ambiguity about the mean and hence the feedback from confidence to economic activity can be also approximated by linearization. In turn, log-linear decision rules facilitate aggregation because the cross-sectional mean becomes a sufficient statistic for tracking aggregate dynamics.

We follow Ilut and Schneider (2014) and solve for the equilibrium law of motion using a guess-and-verify approach: (a) guess the worst case beliefs p^0 ; (b) solve the model assuming that agents have agents have expected utility and beliefs p^0 ; (c) compute the agent's value function V and (d) verify that the guess p^0 indeed achieves the minima.

For step (b) we log-linearize the equilibrium conditions. Details on the recursive representation are in Appendix 6.2. In Appendix 6.3 we present the resulting optimality conditions, which will be a subset of those characterizing the estimated model with additional rigidities introduced in section 4. In what follows we explain the log-linearizing logic by simple expressions for the expected worst-case output at stage 1 (pre-production) and the realized output at stage 2 (post-production). We use the example to illustrate that uncertainty about the firm-level productivity has a first-order effect at the aggregate. We provide a general description of the procedure in Appendix 6.4.

We first find the worst-case steady state by evaluating a deterministic version of the filtering problem and standard first-order conditions under the guessed worst-case belief.²³ Next, we log-linearize the model around the worst-case steady state. To do this, we first log-linearize the expected worst-case output of individual firm l at stage 1:

$$E_t^* \hat{Y}_{l,t}^0 = \hat{A}_t^0 + E_t^* \hat{z}_{l,t}^0 + \hat{Y}_{l,t}^0, \quad (2.24)$$

²³Potential complications arise because the worst-case TFP depends on the level of economic activity. Since the worst-case TFP, in turn, determines the level of economic activity, there could be multiple steady states. We circumvent this multiplicity by treating the posterior variance of the level of idiosyncratic TFP as a parameter and by focusing on the steady state that is implied by that posterior variance.

and the realized output of individual firm l at stage 2:

$$\hat{Y}_{l,t}^0 = \hat{A}_t + \hat{z}_{l,t}^0 + \hat{Y}_{l,t}^0, \quad (2.25)$$

where we use $\hat{x}_t^0 = x_t - \bar{x}^0$ to denote log-deviations from the worst-case steady state and set the trend growth rate γ to zero to ease notation. The worst-case individual output (2.24) is the sum of three components: the current level of aggregate TFP, the worst-case individual TFP, and the input level. The realized individual output (2.25), in turn, is the sum of aggregate TFP, the *realized* individual TFP, and the input level.

We then aggregate the log-linearized individual conditions (2.24) and (2.25) to obtain the cross-sectional mean of worst-case individual output:

$$E_t^* \hat{Y}_t^0 = \hat{A}_t + E_t^* \hat{z}_t^0 + \hat{Y}_t^0, \quad (2.26)$$

and the cross-sectional mean of realized individual output:

$$\hat{Y}_t^0 = \hat{A}_t + \hat{z}_t^0 + \hat{Y}_t^0, \quad (2.27)$$

where we simply eliminate subscript l to denote the cross-sectional mean, i.e., $\hat{x}_t^0 \equiv \int_0^1 \hat{x}_{l,t}^0 dl$.

So far we have characterized the dynamics of output under the worst-case scenario. Our final step is to characterize the dynamics under the true data generating process (DGP). To do this, we feed in the cross-sectional mean of individual TFP, which is constant under the true DGP, into (2.26) and (2.27). Using (2.26), the cross-sectional mean of worst-case output is given by

$$E_t^* \hat{Y}_t = \hat{A}_t + E_t^* \hat{z}_t + \hat{Y}_t, \quad (2.28)$$

where we use $\hat{x}_t = x_t - \bar{x}$ to denote log-deviations from the steady-state under the true DGP. Using (2.27), the realized aggregate output is given by

$$\hat{Y}_t = \hat{A}_t + \hat{Y}_t, \quad (2.29)$$

where we used $\hat{z}_t = 0$ under the true DGP. Importantly, $E_t^* \hat{z}_t$ in (2.29) is not necessarily zero outside the steady state. To see this, combine (2.15) and (2.19) and log-linearize to obtain an expression for $E_t^* \hat{z}_{l,t}$:

$$E_t^* \hat{z}_{l,t} = \varepsilon_{z,z} \hat{z}_{l,t-1|t-1} - \varepsilon_{z,\Sigma} \hat{\Sigma}_{l,t-1|t-1}. \quad (2.30)$$

From (2.10), the posterior variance is negatively related to the level of input \hat{Y} :

$$\hat{\Sigma}_{l,t-1|t-1} = \varepsilon_{\Sigma,\Sigma} \hat{\Sigma}_{l,t-2|t-2} - \varepsilon_{\Sigma,Y} \hat{Y}_{l,t-1}, \quad (2.31)$$

The elasticities $\varepsilon_{z,z}$, $\varepsilon_{z,\Sigma}$, $\varepsilon_{\Sigma,\Sigma}$, and $\varepsilon_{\Sigma,Y}$ are functions of structural parameters and are all positive. We combine (2.30) and (2.31) to obtain

$$E_t^* \hat{z}_{l,t} = \varepsilon_{z,z} \hat{z}_{l,t-1|t-1} - \varepsilon_{z,\Sigma} \varepsilon_{\Sigma,\Sigma} \hat{\Sigma}_{l,t-2|t-2} + \varepsilon_{z,\Sigma} \varepsilon_{\Sigma,Y} \hat{Y}_{l,t-1}. \quad (2.32)$$

Finally, we aggregate (2.32) across all firms:

$$E_t^* \hat{z}_t = -\varepsilon_{z,\Sigma} \varepsilon_{\Sigma,\Sigma} \hat{\Sigma}_{t-2|t-2} + \varepsilon_{z,\Sigma} \varepsilon_{\Sigma,Y} \hat{Y}_{t-1}, \quad (2.33)$$

where we used $\int_0^1 \hat{z}_{l,t-1|t-1} dl = 0$.²⁴

Notice again that the worst-case conditional cross-sectional mean simply aggregates linearly the worst-case conditional mean, $-a_{l,t}$, of each firm. Since the firm-specific worst-case means are a function of idiosyncratic uncertainty, which in turn depend on the firms' scale, equation (2.33) shows that the average level of economic activity, \hat{Y}_{t-1} , has a first-order effect on the cross-sectional average of the worst-case mean. For example, this means that during recessions, firms on average produce less, which leads to lower confidence about their firm-level TFP. This endogenous reduction in confidence further reduces equilibrium hours worked and other economic activity.

3 Propagation mechanism

In this section we characterize the main properties of the propagation mechanism implied by the endogenous firm-level uncertainty. An important part of understanding those dynamics is to explore the way in which the model generates correlated endogenous wedges and co-movement. We find the basic RBC model described in section 2 a particularly useful laboratory for the purpose of facilitating comparison with the standard paradigm. In section 4 we further introduce nominal and real rigidities and estimate the augmented model using a likelihood-based method.

3.1 Co-movement and the labor wedge

Endogenous uncertainty leads to co-movement and a countercyclical labor wedge. This can be analyzed by considering the optimal labor tradeoff of equating the marginal cost to the expected marginal benefit under the worst-case belief E_t^*

$$H_t^\phi = E_t^* (\lambda_t MPL_t) \quad (3.1)$$

In the standard model, there is no expectation on the right-hand side. As emphasized by Barro and King (1984), there hours and consumption move in opposite direction unless there is a TFP or

²⁴This follows from aggregating the log-linearized version of (2.8) and evaluating the equation under the true DGP. Intuitively, since the cross-sectional mean of idiosyncratic TFP is constant, the cross-sectional mean of the Kalman posterior mean estimate is a constant as well.

a labor supply shock (a preference shock to hours worked in agent's utility (2.18)).

Instead, in our model, there can be such co-movement. Suppose that there is a period of low confidence. From the negative wealth effect there is a low consumption, so the standard effect would be to see high labor supply as a result of the high marginal utility of consumption λ_t . However, because the firm chooses hours *as if* productivity is low, there is a counter substitution incentive for hours to be low. To see how the model generates countercyclical labor wedge, note that an increase in ambiguity due to a reduction in labor supply looks, from the perspective of an econometrician, like an increase in the labor income tax.²⁵ The labor wedge can now be easily explained by implicitly defining the labor tax τ_t^H as

$$H_t^\phi = (1 - \tau_t^H)\lambda_t MPL_t$$

Using the optimality condition in (3.1), the labor tax is

$$\tau_t^H = 1 - \frac{E_t^*(\lambda_t MPL_t)}{\lambda_t MPL_t} \quad (3.2)$$

Consider first the linear rational expectations case. There the role of idiosyncratic uncertainty disappears and the labor tax in equation (3.2) is constant and equal to zero. To see this, note our timing assumption that labor is chosen after the aggregate shocks are realized and observed at the beginning of the period. This makes the optimality condition in (3.1) take the usual form of an intratemporal labor decision.²⁶

In our model, the role of idiosyncratic uncertainty does not vanish and instead it shows up in the *as if* expected return to working, formed under the worst-case belief E_t^* . Thus, even if labor is chosen after aggregate shocks are realized and observed, the average idiosyncratic uncertainty has a first-order effect on the cross-sectional average of the worst-case mean, as detailed in section 2.5.

Consider now the econometrician that measures realized hours, consumption and the marginal product of labor as of time t . While agents take the labor decision under E_t^* , the econometrician measures these equilibrium objects under the data generating process which uses the average $\mu = 0$. The difference between the worst-case distribution and the average realization under the econometrician's data generating process produces a labor wedge, which, in log-linear deviations,

²⁵Given the equilibrium confidence process, which determines the worst-case belief E_t^* , the economic reasoning behind the effects of distorted beliefs on labor choice has been well developed by existing work, such as Angeletos and La'O (2009, 2013). There they describe the key income and substitution forces through which correlated higher-order beliefs, a form of confidence shocks, show up as labor wedges in a model where hiring occurs under imperfect information on its return. In addition, Angeletos et al. (2014) emphasize the critical role of beliefs being about the short-run rather than the long-run activity in producing stronger substitution effects. In our setup agents learn about the stationary component of firm-level productivity and therefore the equilibrium worst-case belief typically leads to such stronger substitution effects.

²⁶If we would assume that labor is chosen before the aggregate shocks are realized, there would be a fluctuating labor tax in (3.2) even in the rational expectations model. In that model, the wedge is $\tau_t^H = 1 - \frac{E_{t-1}(\lambda_t MPL_t)}{\lambda_t MPL_t}$, where, by the rational expectations assumptions, E_{t-1} reflects that agents form expectations using the econometrician's data generating process. Crucially, in such a model, the labor wedge τ_t^H will not be predictable using information at time $t-1$, including the labor choice, such that $E_{t-1}\tau_t^H = 0$. In contrast, our model with learning produces predictable, countercyclical, labor wedges.

is inversely proportional to the time-varying confidence.

In a period of low confidence, the ratio between the expected benefit to working under the worst-case belief compared to the econometrician’s measure of $\lambda_t MPL_t$ is typically lower. Thus, the econometrician rationalizes the ‘surprisingly low’ labor supply by a high labor tax τ_t^H . In turn, the low confidence is generated endogenously from a low level of average economic activity, as reflected in the lower cross-sectional average of the worst-case mean, as given by equation (2.33). Therefore, the econometrician finds a systematic negative relationship between economic activity and the labor income tax. This relationship is consistent with empirical studies that suggest that in recessions labor falls by more than what can be explained by the marginal rate of substitution between labor and consumption and the measured marginal product of labor (see for example Shimer (2009) and Chari et al. (2007)).

Finally, for an ease of exposition, we have described here the behavior of the labor wedge by ignoring the potential effect of experimentation on the optimal labor choice. This effect adds an additional reason why labor moves ‘excessively’, from the perspective of an observer that only uses equation (3.1) to understand labor movements. As discussed later in section 3.4, experimentation amplifies the effects of uncertainty during the short-run, and thus leads to even more variable labor wedges, while it dampens fluctuations in the medium-run.

3.2 Intertemporal discount factor wedge

Uncertainty also affects the consumption-savings decision of the household. This is reflected in the Euler condition for the risk-free asset:

$$1 = \beta R_t E_t^* (\lambda_{t+1} / \lambda_t) \tag{3.3}$$

As with the labor wedge, let us implicitly define an intertemporal savings wedge:

$$1 = (1 + \tau_t^B) \beta R_t E_t (\lambda_{t+1} / \lambda_t)$$

Importantly, this wedge is time varying, since the bond is priced under the uncertainty adjusted distribution, E_t^* , which differs from the econometrician’s DGP, given by E_t . By substituting the optimality condition for the interest rate from (3.3), the wedge becomes:

$$1 + \tau_t^B = \frac{E_t^* \lambda_{t+1}}{E_t \lambda_{t+1}} \tag{3.4}$$

Equation (3.4) makes transparent the predictable nature of the wedge. In particular, during low confidence times, the representative household acts *as if* future marginal utility is high. This heightened concern about future resources drives up demand for safe assets and leads to a low interest rate R_t . However, from the perspective of the econometrician, the measured average

marginal utility at $t + 1$ is not particularly high. To rationalize the low interest rate without observing large changes in the growth rates of marginal utility, the econometrician recovers a high savings wedge τ_t^B .

Therefore, the model offers a mechanism to generate movements in the relevant discount factor that arise endogenously as a countercyclical desire to save in risk-free assets. On the one hand, the endogenous wedge is not simply a by-product, but it helps understand the model mechanism. Indeed, even a supply shock such as an aggregate TFP shock will endogenously generate features resembling a demand shock, such as those arising from an intertemporal preference shock. On the other hand, the endogenous wedge is a type of overidentifying restriction that is consistent with the data, without requiring exogenous shocks to the demand for safety, a source of stochasticity used in recent macroeconomic models, such as Fisher (2015) or Christiano et al. (2015).

3.3 Excess return

A similar logic applies to the countercyclical excess return. The Euler condition for capital is

$$\lambda_t = \beta E_t^*[\lambda_{t+1} R_{t+1}^K]$$

Under our linearized solution, using equation (3.3), we get $E_t^* R_{t+1}^K = R_t$, where $E_t^* R_{t+1}^K$ is the expected return on capital under the worst-case belief. As with the intertemporal savings wedge, let us define the measured excess return wedge as

$$E_t R_{t+1}^K = R_t(1 + \tau_t^K)$$

As with bond pricing, this wedge is time-varying and takes the form

$$1 + \tau_t^K = \frac{E_t R_{t+1}^K}{E_t^* R_{t+1}^K} \tag{3.5}$$

During low confidence times, demand for capital is ‘surprisingly low’. This is rationalized by the econometrician, measuring R_{t+1}^K under the true DGP, as a high ex-post excess return $R_{t+1}^K - R_t$, or as a high wedge τ_t^K in equation (3.5). In the linearized solution, the excess return, similarly to the labor tax and the discount factor wedge, is inversely proportional to the time-varying confidence. In times of low economic activity, when confidence is low, the measured excess return is high.

Thus, the model can generate a countercyclical labor and savings wedge at the same time when the measured premia on uncertain assets is high.²⁷ This arises from any type of shock that moves the economic activity.

²⁷The pricing logic can be extended to defaultable corporate bonds. This will likely generate countercyclical excess bond premia, as documented for example by Gilchrist and Zakrajšek (2012).

Table 1: Parameters for the impulse response exercise

γ	α	β	ϕ	δ	η	$\bar{\Sigma}$	\bar{g}	ρ_z	σ_z	ρ_A	ρ_Δ
1.004	0.3	0.99	0	0.025	2	0.1	0.2	0.5	0.4	0.95	0.95

3.4 Impulse response analysis

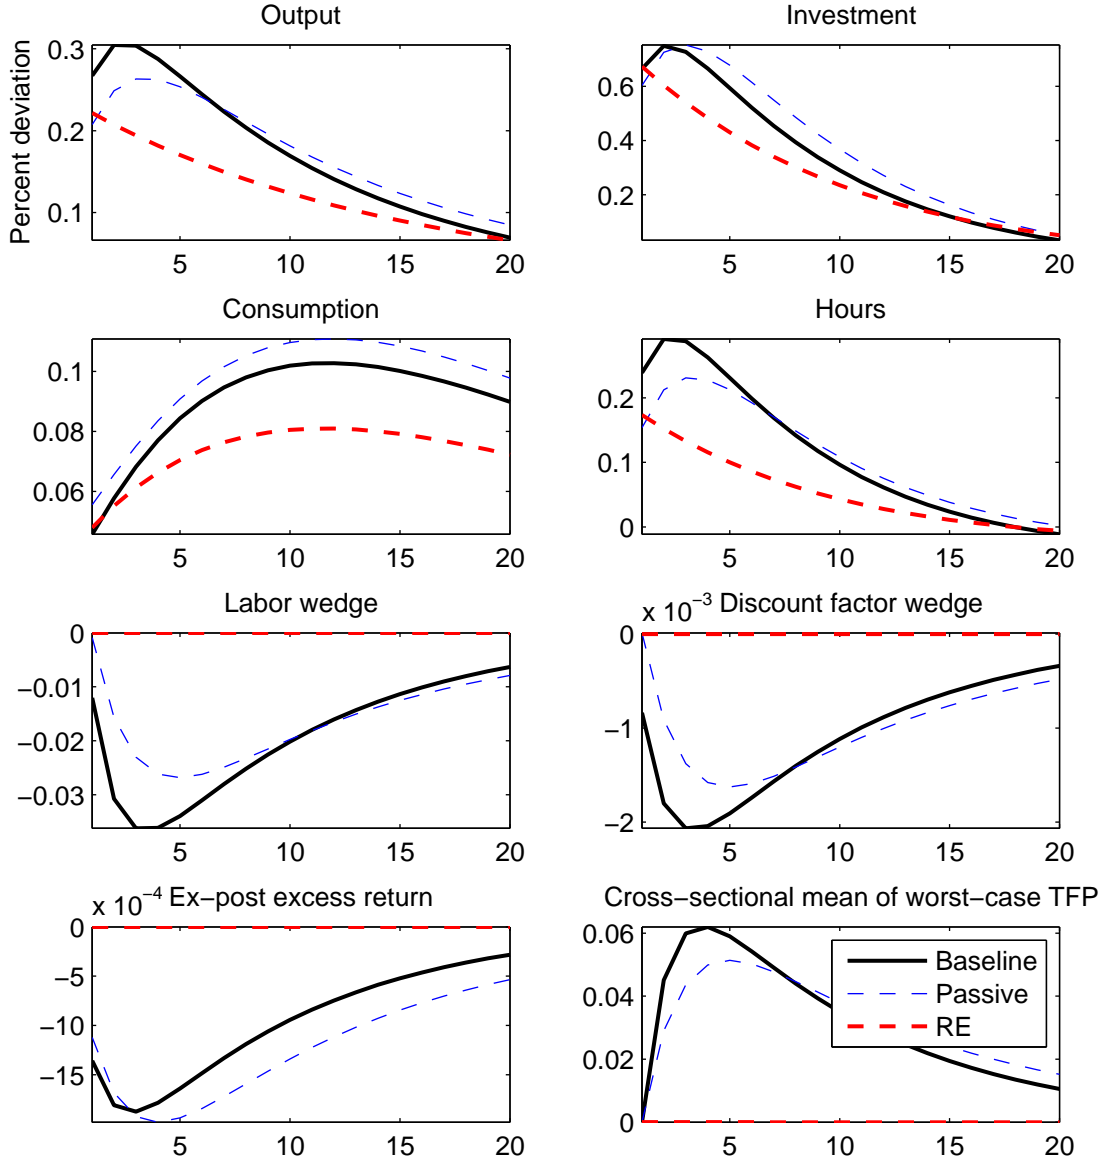
In the previous sections we have described some of the analytics behind the emergence of counter-cyclical wedges. In this subsection we illustrate the dynamics through impulse response functions. In order to facilitate comparison with a standard RBC model, we use common values from the literature whenever possible and explain the role of the other parameters. Table 1 summarizes the parameters used in our exercise. These parametrization serves an illustration purpose, while in section 6 we estimate a model rich in various rigidities.

The magnitude of the feedback loop between uncertainty and economic activity is determined by three factors. The first factor is the variability of labor, controlled by the elasticity of labor supply. For illustration purposes, we set the inverse Frisch elasticity ϕ equal to zero following the indivisible labor model by Hansen (1985) and Rogerson (1988). Second, the parameters that are related to the idiosyncratic processes control how changes in inputs translate to changes in the posterior variance. We choose $\rho_z = 0.5$ and $\sigma_z = 0.4$ for the idiosyncratic TFP process. Idiosyncratic TFP is less persistent than the aggregate, which is in line with the finding in Kehrig (2015). The values imply a cross-sectional standard deviation of TFP of 0.46 and is in line with the estimates found in Bloom et al. (2014) using the establishment-level data. David et al. (2015) estimate the posterior variance of a firm-specific shock (in the context of our model, a TFP shock) to be around 8–13%. We choose the worst-case steady state posterior variance so that at the zero-risk steady state the posterior variance $\bar{\Sigma}$ is 10%.²⁸ Finally, the size of the entropy constraint η determines how changes in the posterior standard deviation translate into changes in confidence. Ilut and Schneider (2014) argue that a reasonable upper bound for η is 2, based on the view that agents’ ambiguity should not be “too large”, in a statistical sense, compared to the variability of the data.

Figure 1 plots the impulse response to a positive TFP shock. In addition to the response from the baseline model (labeled ‘Baseline’), we also report the responses from the model with passive learning (labeled ‘Passive’), in which the firm does not internalize the effect of its input choice on its future uncertainty, and the standard rational expectation (RE) RBC model (labeled ‘RE’). The solution to the RE model is obtained by simply setting the entropy constraint η to zero. In this

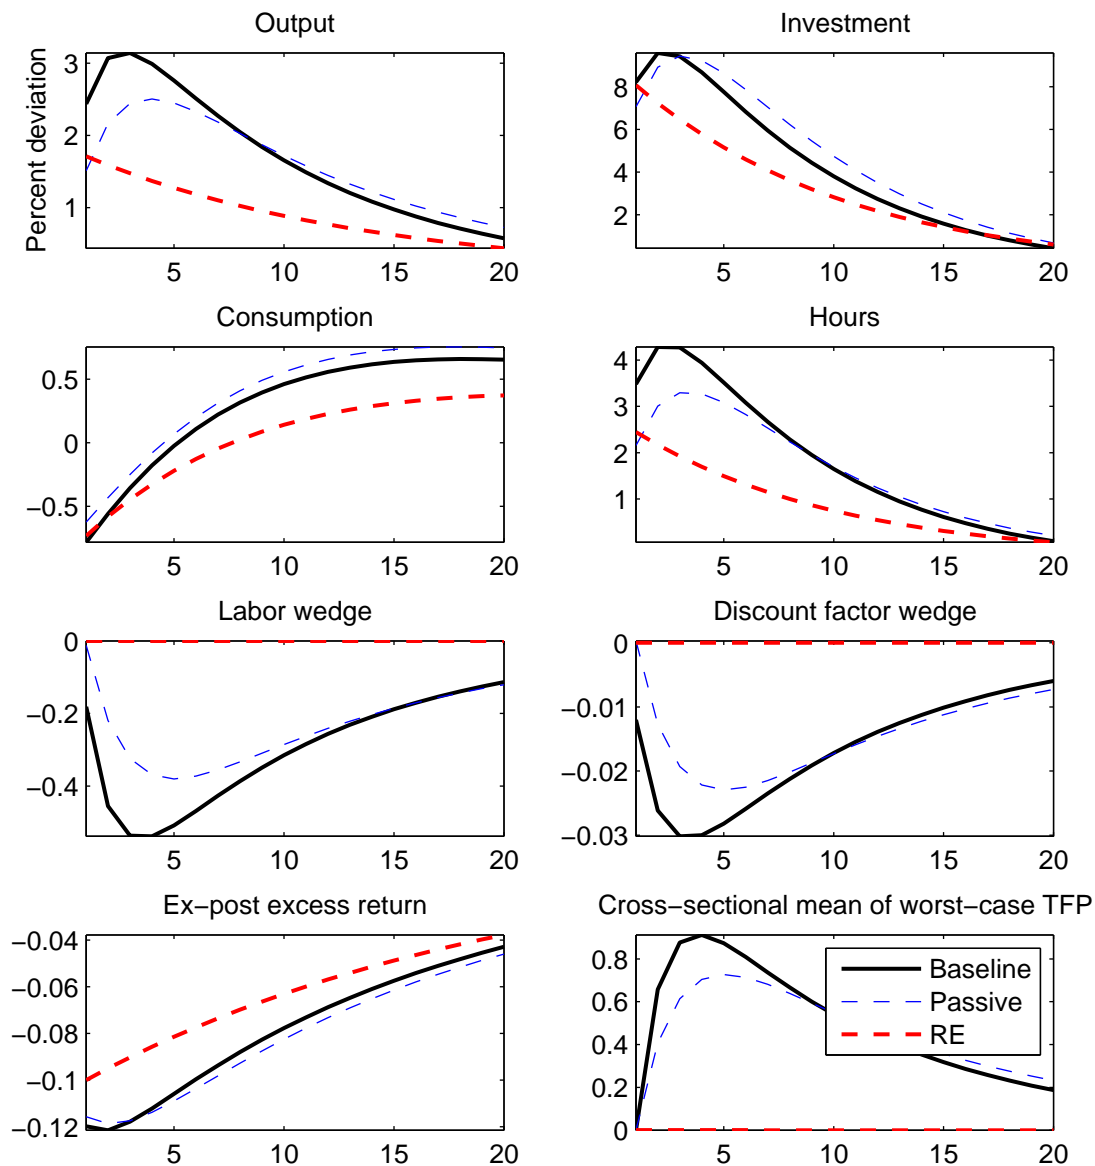
²⁸We re-parameterize the model so that we take the worst-case steady state posterior variance $\bar{\Sigma}^0$ of idiosyncratic TFP as a parameter. This posterior variance, together with ρ_z and σ_z , will pin down the standard deviation of the luck component σ_ν . The zero-risk steady state is the ergodic steady state of the economy where optimality conditions take into account uncertainty and the data is generated under the econometrician’s DGP. Appendix 6.4 provides additional details. Under the resulting parametrization, the Kalman gain at the zero-risk steady state, normalizing the level of input to one, is 0.47. This implies that an observation from four quarters ago will receive a weight $(1 - 0.47)^4 \approx 0.08$.

Figure 1: Impulse response to an aggregate TFP shock



Notes: Thick black solid line ('Baseline') is our baseline model with active learning, thin blue dashed line ('Passive') is the model with passive learning, and thick red dashed line ('RE') is the frictionless, rational expectations, RBC model.

Figure 2: Impulse response to a financial shock



Notes: Thick black solid line ('Baseline') is our baseline model with active learning, thin blue dashed line ('Passive') is the model with passive learning, and thick red dashed line ('RE') is the frictionless, rational expectations, RBC model.

case, agents think in terms of single probabilities and the model reduces to a rational expectation model. Note that when $\eta = 0$, firm-level learning cancels out in the aggregate due to linearization and the law of large numbers.

Compared to the RE version, our model generates amplified and hump-shaped response in output, investment, and hours. These dynamics are due to the endogenous variation in firms' confidence. In response to a positive TFP shock, firms (on average) increase their inputs, such as hours and the capital utilization rate. The increase in inputs lowers uncertainty which implies that firms contemplate a narrower set of conditional probabilities; the worst-case scenarios become less worse. As a result, the agent acts *as if* the mean idiosyncratic productivities are higher and this may further stimulate economic activity. At the same time, from the econometrician's perspective, the labor supply, the return on risk-free bonds and the demand for capital are surprisingly high. Thus, various wedges co-move: the labor wedge, the discount factor wedge and the ex-post excess return on capital decline.

Finally, we compare our baseline impulse response with the response from the passive learning model. Initially the output and hours responses of the baseline model with active learning are larger than those of passive learning. In the medium run, however, the responses of passive learning become larger. This is due to a dynamic interaction of two opposing forces. On one hand, higher production during booms increases the value of experimentation because it raises the marginal benefit of an increase in the expected worst-case technology. On the other hand, there is an offsetting effect coming from a reduction in posterior variance. Since the level of posterior variance is downward convex in the level of inputs, the marginal reduction in posterior variance due to an increase in inputs is smaller during booms. During the initial period of a positive technology shock, the first effect dominates the second. As the economy slows down, the second effect becomes more important.

Figure 2 shows the impulse response to a financial shock. In a standard RBC model the increase in expected return to capital, as a result of the reduced financial wedge shown in the last row of the first column, produces a standard intertemporal and wealth effect on consumption. The former tends to be stronger since the financial shock affects directly the expected return to investment while the former is weaker since the shock does not directly impact the resource constraint. As a consequence, investment rises through a fall in consumption. In turn, the fall in consumption stimulates hours worked as marginal utility of consumption is higher. Consumption and hours co-move positively only in the long run after the positive productivity effect of capital accumulation kicks in. In our model, the increase in hours raises average confidence and further stimulates economic activity. This leads to a stronger positive wealth effect, as well as a higher perceived marginal return to labor. The positive effect can be strong enough to make consumption increase and the labor wedge can be low enough to simultaneously lead to high equilibrium labor. Therefore, our model can generate positive co-movement in consumption and labor after a demand shock much quicker than the typical effects present in the RE model that arise from endogenous capital accumulation. At the same time, the

financial shock, similarly to the technology shock, manifests as correlated countercyclical wedges.

4 Quantitative analysis

We now estimate the model by conducting Bayesian inference. We introduce variable capital utilization and embed real and nominal rigidities widely used in the literature. These extensions allow us to compare our learning mechanism with traditional frictions and to investigate the extent to which those frictions, in turn, are replaceable by our mechanism.

4.1 Bayesian estimation

We start by describing the additional features that we introduce to the estimated model. The production function with capital utilization is

$$\tilde{Y}_{l,t} = (U_{l,t}K_{l,t-1})^\alpha (\gamma^t H_{l,t})^{1-\alpha}$$

and $a(U_{l,t})K_{l,t-1}$ is an utilization cost that reduces dividends in equation (2.7).²⁹ We modify the representative household's utility (2.18) to allow for habit persistence in consumption:

$$U_t(C; s^t) = \ln(C_t - bC_{t-1}) - \frac{H_t^{1+\phi}}{1+\phi} + \beta \min_{\mu_{l,t} \in [-a_{l,t}, a_{l,t}], \forall l} E^\mu [U_{t+1}(C; s^t, s_{t+1})],$$

where $b > 0$ is a parameter. We also introduce a standard investment adjustment cost:

$$K_{l,t} = (1 - \delta)K_{l,t-1} + \left\{ 1 - \frac{\kappa}{2} \left(\frac{I_{l,t}}{I_{l,t-1}} - \gamma \right)^2 \right\} I_{l,t}, \quad (4.1)$$

where $\kappa > 0$ is a parameter. We introduce nominal rigidity by considering Calvo-type price and wage stickiness along with monopolistic competition.³⁰ We assume a Taylor-type reaction function:

$$\frac{R_t}{\bar{R}} = \left(\frac{R_{t-1}}{\bar{R}} \right)^{\rho_R} \left[\left(\frac{\pi_t}{\bar{\pi}} \right)^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}} \right)^{\phi_Y} \right]^{1-\rho_R} \epsilon_{R,t}, \quad \epsilon_{R,t} \sim N(0, \sigma_R^2),$$

²⁹We specify: $a(U) = 0.5\chi_1\chi_2U^2 + \chi_2(1 - \chi_1)U + \chi_2(0.5\chi_1 - 1)$, where χ_1 and χ_2 are parameters. We set χ_2 so that the steady-state utilization rate is one. The cost $a(U)$ is increasing in utilization and χ_1 determines the degree of the convexity of utilization costs. In a linearized equilibrium, the dynamics are controlled by the χ_1 .

³⁰To avoid complications arising from directly embedding infrequent price adjustment into firms, we follow Bernanke et al. (1999) and assume that the monopolistic competition happens at the ‘‘retail’’ level. Retailers purchase output from firms in a perfectly competitive market, differentiate them, and sell them to final-goods producers, who aggregate retail goods using the conventional CES aggregator. The retailers are subject to the Calvo friction and thus can adjust their prices in a given period with probability $1 - \xi_p$. To introduce sticky wages, we assume that households supply differentiated labor services to the labor packer with a CES technology who sells the aggregated labor service to firms. Households can only adjust their wages in a given period with probability $1 - \xi_w$.

where $\bar{\pi}$ is the inflation target, $\bar{R} = \bar{\pi}\gamma/\beta$, and \bar{Y} is output along the balanced growth path. ρ_R , ϕ_π , and ϕ_Y are parameters and $\epsilon_{R,t}$ is a monetary policy shock.

The stochastic structure of our economy is the same as in the baseline model of section 3. There are three sources of aggregate shocks: the technology shock A_t , the financial shock Δ_t^k , and the monetary policy shock $\epsilon_{R,t}$. The technology and financial shocks follow the AR(1) processes in (2.3) and (2.14), respectively.

Of particular importance for the identification discussed below, note that combining the Euler equation for the risk-free asset, $1 = E_t^* M_{t+1} R_t$, with that for capital accumulation and rearranging, we obtain $\Delta_t^k \simeq E_t^* R_{t+1}^k - R_t$. Therefore, the unobserved financial shock is tightly connected to a measure of excess return on capital, for which we will use an observable in our estimation.

Because the aggregate law of motion is linear, we can use standard Bayesian techniques as described in An and Schorfheide (2007) to estimate the model. The sample period is 1985Q1–2014Q4. The data is described in Appendix 6.5. The vector of observables is

$$[\Delta \ln Y_t, \Delta \ln H_t, \Delta \ln I_t, \Delta \ln C_t, \ln \pi_t, \ln R_t, Spread_t].$$

where $Spread_t$ is the Baa corporate bond yield relative to the yield of Treasury bond with ten-year maturity.³¹ We assume that the model counterpart of $Spread_t$ is the excess-return on capital so that $Spread_t = R_t^k - R_{t-1}$.³²

We fix a small number of parameters prior to estimation. The growth rate of technology γ , the depreciation rate of capital δ , and the share of government spending to output \bar{g} are set to the values in Table 1. The prior distributions for other structural parameters are collected in Tables 2 and 3. We reparametrize the parameter that determines the size of ambiguity (η) and instead estimate 0.5η . We set a Beta prior for 0.5η , so that the lowest value corresponds to rational expectations and the highest value corresponds to the upper bound $\eta = 2$ suggested by Ilut and Schneider (2014). Since all our estimation exercises have less structural shocks than observables, we add i.i.d. measurement error to observables except for the interest rate and the financial spread. At the posterior mean, measurement error explains 1 percent of variation of a particular observable. We place discipline on the estimation and set relatively small and tight priors for the measurement error so that the model tries to explain observables through movements in endogenous variables as much as possible. To facilitate comparison we also estimate the RE version of the model.

³¹We also used the spread constructed in Gilchrist and Zakrajšek (2012) and obtained similar results.

³²To understand what causes variations in the spread, it is useful to decompose it into three components:

$$Spread_t = (E_{t-1}^* R_t^k - R_{t-1}) + (R_t^k - E_t' R_t^k) + (E_t' R_t^k - E_{t-1}^* R_t^k), \quad (4.2)$$

where we use E_t' to use worst-case expectations at the end of stage 1 (after the realization of period t aggregate shock but before the realization of idiosyncratic shocks). The financial shock Δ_t^k causes variations in the first component, changes in confidence cause variations in the second component, and innovations to aggregate shocks (such as TFP shocks) cause variations in the third component.

Table 2: Estimated parameters: preference, technology, and policy

		Prior			Posterior mode	
		Type	Mean	Std	Ambiguity	RE
$100(\gamma - 1)$	Growth rate	N	0.5	0.3	0.57	0.41
					[0.56, 0.58]	[0.35, 0.46]
$100(\beta^{-1} - 1)$	Discount factor	G	0.1	0.05	0.06	0.05
					[0.03, 0.08]	[0.02, 0.10]
$100(\bar{\pi} - 1)$	Net inflation	N	0.2	0.05	0.39	0.23
					[0.34, 0.40]	[0.12, 0.29]
α	Capital share	B	0.3	0.04	0.41	0.26
					[0.40, 0.42]	[0.21, 0.34]
ϕ	Inverse Frisch elasticity	G	1	0.4	0.20	1.14
					[0.17, 0.21]	[1.13, 1.51]
χ_1	Utilization cost	G	0.7	0.3	0.52	0.51
					[0.49, 0.54]	[0.40, 0.83]
b	Consumption habit	B	0.3	0.1	0.57	0.95
					[0.54, 0.59]	[0.92, 0.97]
κ	Investment adj. cost	G	1.5	0.5	0.09	4.77
					[0.08, 0.09]	[4.52, 4.93]
ξ_p	Calvo price	B	0.5	0.2	0.44	0.89
					[0.42, 0.45]	[0.85, 0.93]
ξ_w	Calvo wage	B	0.5	0.2	0.02	0.82
					[0.00, 0.02]	[0.57, 0.93]
ρ_R	Interest smoothing	B	0.7	0.1	0.66	0.40
					[0.63, 0.66]	[0.31, 0.46]
ϕ_π	Inflation response	N	1.5	0.05	1.65	1.61
					[1.62, 1.66]	[1.55, 1.71]
ϕ_Y	Output response	N	0.2	0.05	0.05	0.25
					[0.05, 0.06]	[0.19, 0.35]
ρ_z	Idiosyncratic TFP	B	0.5	0.05	0.44	–
					[0.43, 0.45]	
σ_z	Idiosyncratic TFP	G	0.4	0.1	1.13	–
					[1.13, 1.15]	
0.5η	Entropy constraint	B	0.5	0.2	0.85	0
					[0.81, 0.86]	
$\bar{\Sigma}$	SS posterior variance	G	0.1	0.05	0.17	–
					[0.09, 0.31]	

Notes: ‘Ambiguity’ corresponds to our baseline model with endogenous uncertainty and ‘RE’ corresponds to its rational expectations version. *B* refers to the Beta distribution, *N* to the Normal distribution, *G* to the Gamma distribution, *IG* to the Inverse-gamma distribution. 95% posterior intervals are in brackets and are obtained from draws using the random-walk Metropolis-Hasting algorithm.

Table 3: Estimated parameters: shocks and measurement errors

		Prior			Posterior mode	
		Type	Mean	Std	Ambiguity	RE
ρ_A	Aggregate TFP	B	0.9	0.05	0.93	0.97
					[0.91, 0.94]	[0.90, 0.96]
ρ_Δ	Financial	B	0.9	0.05	0.97	0.95
					[0.97, 0.98]	[0.93, 0.97]
$100\sigma_A$	Aggregate TFP	IG	1	1	0.28	0.48
					[0.24, 0.29]	[0.34, 0.61]
$100\sigma_\Delta$	Financial	IG	0.1	1	0.04	0.02
					[0.03, 0.05]	[0.02, 0.03]
$100\sigma_R$	Monetary policy	IG	0.1	1	0.06	0.05
					[0.06, 0.09]	[0.03, 0.08]
$100\sigma_Y$	Meas. error output	G	0.060	0.03	0.37	0.56
					[0.38, 0.40]	[0.50, 0.62]
$100\sigma_H$	Meas. error hours	G	0.073	0.03	0.44	0.54
					[0.43, 0.47]	[0.48, 0.60]
$100\sigma_I$	Meas. error investm.	G	0.18	0.08	1.12	1.62
					[1.08, 1.13]	[1.48, 1.75]
$100\sigma_C$	Meas. error consum.	G	0.050	0.02	0.32	0.44
					[0.30, 0.35]	[0.40, 0.49]
$100\sigma_\pi$	Meas. error inflation	G	0.024	0.01	0.21	0.36
					[0.20, 0.24]	[0.33, 0.40]
Log marginal likelihood					3262	3096

Notes: See notes for Table 2.

4.2 Results

The posterior estimates are collected in Tables 2 and 3. Our model with endogenous uncertainty is labeled ‘Ambiguity’ and its rational expectations version is labeled ‘RE’.

We first study the rational expectations (RE) version of the model. The row labeled ‘RE’ in Table 4 reports the contribution of a financial shock for business cycles. The Table shows that the financial shock accounts for a modest fraction of variables such as output and hours. For example, the financial shock explains only 19 and 4 percents of variations in output and hours, respectively. This implies that in the RE model the TFP shock is the major driver of business cycles.³³ Next, the column labeled ‘RE’ in Table 2 shows that the estimated degree of rigidities is high. This happens because the model tries to be consistent with the observed data by generating persistent dynamics and co-movement. Nevertheless, as shown in Figures 9 and 10 in Appendix 6.6, neither the TFP nor the financial shock end up generating that co-movement. As is common in New Keynesian models, hours fall after an improvement in technology. In response to a reduction in financial wedge, output, investment, and hours increase but consumption slightly declines.

³³The contribution of a monetary policy shock is negligible.

Table 4: Theoretical variance explained by financial shock

	Output	Hours	Investm.	Consum.	Inflation	Interest
Ambiguity	0.56	0.61	0.61	0.47	0.67	0.72
RE	0.19	0.04	0.34	0.69	0.00	0.01

Notes: ‘Ambiguity’ corresponds to our baseline model with endogenous uncertainty and ‘RE’ corresponds to its rational expectations version. We report the percent of variance at the business cycle frequency (6–32 quarters) that can be explained by the financial shock.

As a result, the RE model fails to explain aggregate fluctuations in the data. To see this, we report the summary statistics of our observables and their model counterparts at the posterior mode, computed by the Kalman smoothing algorithm in the panel labeled ‘RE’ in Table 5. Figure 7 in Appendix 6.6 visualizes this by plotting the historical path of observables along with the model counterparts. Due to the high estimated habit and adjustment cost, the model understates the standard deviations for all endogenous variables, except for inflation. The predicted output, investment, and consumption are negatively correlated with data. To understand the discrepancy between the model and the data, first consider the role of the financial shock. Recall that in the estimation the financial shock is tightly connected to the observed credit spread. Since empirically the spread increases during recessions, this implies that the RE model predicts a fall in output, investment, and hours but an increase in consumption during recessions. Next, consider the role of the TFP shock. Since this shock explains most of variations in hours (95 percent), the model predicts an increase in TFP during recessions, which makes hours fall but output, investment, and consumption increase. The combination of those two shocks make model-implied hours positively correlated with data and model-implied output, investment, and especially consumption negatively correlated with the data (panel ‘RE’ in Table 5). In addition, the model cannot replicate the co-movement patterns in the data: for example, in the data the correlation between hours and consumption growths is 0.52 while in the model it is negatively correlated at -0.44.

Next, consider our model with endogenous uncertainty. The column labeled ‘Ambiguity’ in Table 2 shows that the estimated degree of rigidities is substantially smaller compared to the RE version. For example, the estimated habit parameter is 40% smaller and the investment adjustment cost is negligible. The prices are adjusted every $1/(1 - 0.44) \approx 2$ quarters and wages are adjusted roughly every quarter, instead of the 9 and 6 quarters, respectively, in the RE model. This is because our model can generate co-movement and persistent dynamics without relying on traditional rigidities.

To illustrate the ability to produce co-movement, Figure 3 shows the estimated impulse response to a financial shock. In addition to the baseline impulse response (labeled ‘Baseline’), we also display the response for a version of the model in which we set $\eta = 0$ and freeze the remaining parameters

Table 5: Fit of observables

Model	Statistic	Output	Hours	Investm.	Consum.	Inflation
	$\sigma(\text{data})$	0.60	0.73	1.79	0.50	0.24
Ambiguity	$\sigma(\text{model})$	0.64	0.47	1.29	0.35	0.22
	$\rho(\text{data, model})$	0.84	0.76	0.77	0.74	0.48
RE	$\sigma(\text{model})$	0.11	0.37	0.47	0.03	0.53
	$\rho(\text{data, model})$	-0.16	0.60	-0.05	-0.30	0.49

Notes: ‘Ambiguity’ refers to our baseline model with endogenous uncertainty and ‘RE’ refers to its rational expectations version. $\sigma(\text{data})$ and $\sigma(\text{model})$ refer to standard deviations of the observables and the model counterparts expressed in percentage points. $\rho(\text{data, model})$ is the correlation between data and model.

to the baseline estimated values (labeled ‘RE counterfactual’).³⁴ In the RE counterfactual, an expansionary financial shock leads to a transitory increase in output, hours, and investment but a decline in consumption. In contrast, in our model, output, hours, investment, and consumption all increase in a persistent and delayed manner, with peak responses around 20 quarters after the shock. Because of this, the financial shock emerges as a major driver of the business cycle in our model (the row labeled ‘Ambiguity’ in Table 4). For example, the financial shock explains 56 and 61 percent of output and hours variances, compared to the 19 and 4 percent in the RE version, respectively. That our learning mechanism changes inference about the source of the business cycle and makes demand shock important is one of the main results of this paper.

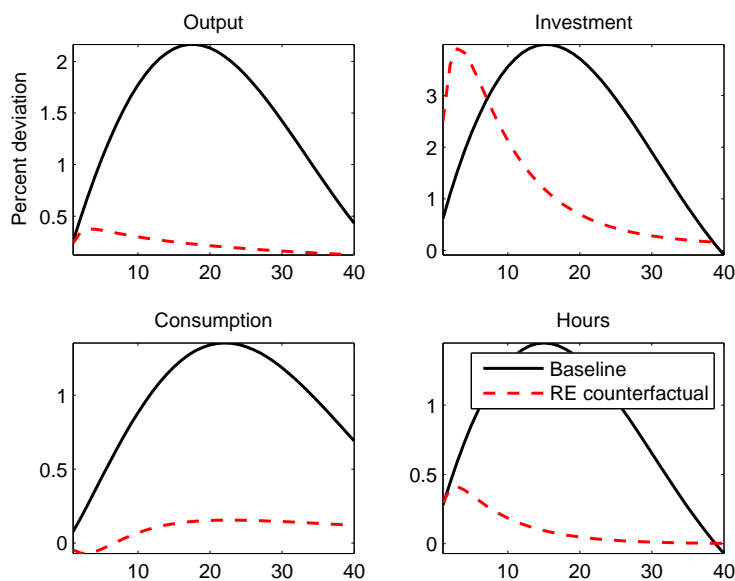
Endogenous uncertainty improves the model performance compared to the RE version in terms of the fit of observables (Panel ‘Ambiguity’ in Table 5). In Figure 8 in Appendix 6.6, the smoothed variables from the model tracks the path of observables quite well. In addition, our model (correctly) predicts that consumption and hours growth are positively correlated (0.72 versus 0.52 in the data). The model fit, measured in terms of log marginal likelihood computed by the Geweke’s modified harmonic mean estimator, which penalizes more parameters, is improved as well. Table 3 reports that in our model the marginal likelihood is 166 log points larger than the RE version.

An additional cross-equation restriction implied by our friction is that part of the movement in the credit spread can be explained by endogenous variation in confidence. Figure 4 visualizes this result by plotting the spread and the model counterpart computed from the Kalman smoother. Note that because there is no measurement error in the spread, the model and data spread is identical. The endogenous variation in the spread is labeled ‘Model (ambiguity component)’.³⁵ Since agents lose confidence during recessions, the movement in confidence can explain a substantial fraction of

³⁴In other words, all parameters except for η are set to the values reported under the column ‘Ambiguity’ in Table 2 and 3.

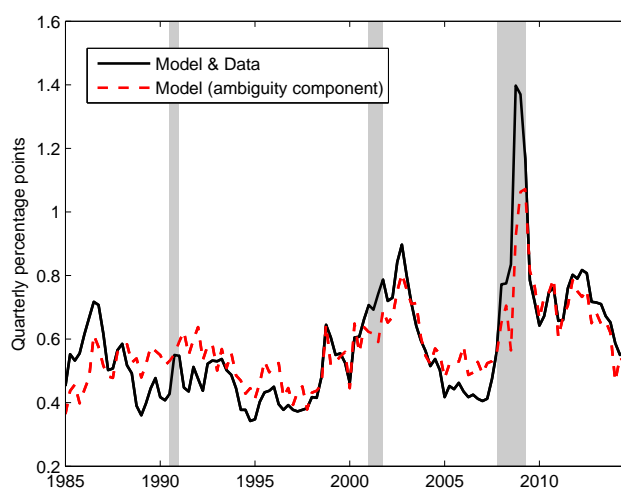
³⁵This endogenous variation corresponds to the second component in the spread decomposition of equation (4.2).

Figure 3: Estimated impulse response to a financial shock




Notes: Black solid line ('Baseline') is the impulse response from our baseline model with endogenous uncertainty and the red dashed line ('RE counterfactual') is the response where we set the entropy constraint $\eta = 0$ while holding other parameters at their estimated values.

Figure 4: Credit spread



Notes: We plot the Baa corporate bond spread along with the model counterpart computed from the Kalman smoother. Black solid line ('Model & Data') is the data and the model, red dashed line ('Model (ambiguity component)') is the spread generated by endogenous variation in confidence.

Figure 5: Time series of firm-level uncertainty



Notes: Black dashed line (‘Data’) is the dispersion of firm-level capital return forecasts across analysts constructed by Senga (2015) and red solid line (‘Model’) is the model counterpart computed from the Kalman smoother.

the variation in the spread, which is countercyclical.

The results point to an important takeaway on the role of financial shocks. On the one hand, as shown by Figure 4, part of the observed spread is explained by the endogenous model-implied component and therefore there is less empirical need for finding financial shocks. On the other hand, the propagation mechanism produced by our information friction makes the estimated shocks, whose time-series properties are disciplined by the observable proxy, have large and empirically plausible effects on the observed macroeconomic aggregates.

4.3 Time-series evidence from firm-level survey data

While our theory puts tight restrictions on time-varying uncertainty through the feedback between learning and economic activity, we can use evidence outside the model to test its predictions. In this section we show that our model-implied firm-level confidence process lines up well with the time-series pattern in uncertainty directly measured from the micro survey data.

Our measure of confidence is the cross-sectional average dispersion of firm-level capital return forecasts. We use a series constructed by Senga (2015) using I/B/E/S and Compustat data.³⁶ For each firm, Senga (2015) measures the coefficient of variation across analysts’ forecasts of the return on capital for that firm. Taking the cross-sectional average across firms of that coefficient of variation results in a time-series measure. For the model counterpart, we calculate the coefficient of variation of capital return forecasts by computing expected capital returns implied by the set of

³⁶We thank Tatsuro Senga for generously sharing his data.

Table 6: Changing the Taylor rule coefficient on the credit spread

ϕ_{spread}	Std. of output growth	
	Baseline	Fixed uncertainty
0	0.64	0.64
-0.5	0.62	0.63
-0.8	0.59	0.63
-1	0.56	0.63

Notes: The standard deviation of output growth as the Taylor rule coefficient on the credit spread ϕ_{spread} is changed from the original value ($\phi_{spread} = 0$). All other parameter values are fixed to the estimated values. ‘Baseline’ refers to our baseline model with endogenous uncertainty. ‘Fixed uncertainty’ refers to a counterfactual economy where the path of uncertainty is fixed exogenously to the original path computed from the economy with $\phi_{spread} = 0$.

productivity process (2.15).³⁷ To control for low-frequency variations in the data possibly stemming from the unbalanced panel and other factors, we detrend both the data and the model using a linear trend. Figure 5 reports Senga (2015)’s uncertainty measure along with the model counterpart. The model tracks the data reasonably well and the increases in uncertainty around the 2000-2001 and 2007-2009 recessions are of similar magnitudes.

To conclude, the time-series variation in uncertainty implied from our estimated model broadly matches firm-level data that was not used in the estimation. This external validation provides additional evidence that the our endogenous uncertainty mechanism is empirically plausible.

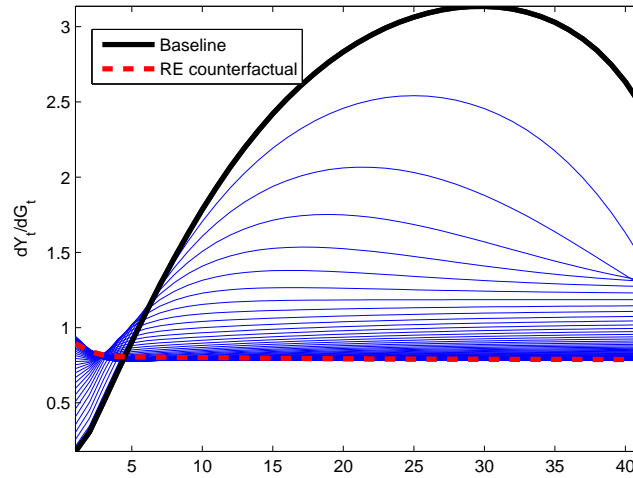
4.4 Policy implications of endogenous uncertainty

Finally, the fact that in our model uncertainty is endogenous has important policy implications. To illustrate this point, we conduct two policy experiments. First, we evaluate the impact of modifying the Taylor rule to incorporate an adjustment to the credit spread. Table 6 reports the standard deviation of output growth as we keep all parameters at their baseline estimated values but change the Taylor rule coefficient on the credit spread ϕ_{spread} from the original value of zero in the estimated model. The standard deviation decreases as monetary policy responds more aggressively to the spread movements. For example, the standard deviation decreases from 0.64 to 0.56 when the coefficient ϕ_{spread} decreases from 0 to -1 .

We find that the reduction in output variability in this counterfactual policy intervention comes

³⁷In terms of mapping the model to data, the idea here is that the representative agent samples experts’ forecast and aggregates them when making decisions. Since the agent is ambiguity averse, stronger disagreement among experts about conditional firm-level mean returns generates lower confidence in probability assessments of the future.

Figure 6: Government spending multiplier



Notes: We plot the government spending multiplier for output. We hit the economy with a positive government spending shock at $t = 1$ and the path of government spending follows an AR(1) process. Thick black solid line (‘Baseline’) is our model with endogenous uncertainty and thick red dashed line (‘RE counterfactual’) is the multiplier where we set the entropy constraint $\eta = 0$ while holding other parameters at their estimated values. In thin blue solid lines, we plot the multipliers for various values of the entropy constraint in an increment of 0.02 from $\eta = 0$ to the estimated baseline value of $\eta = 1.7$.

from stabilizing the endogenous variation in uncertainty. To see this, we also show the effects of policy changes in the counterfactual economy where the path of uncertainty is fixed to the original one (labeled “Fixed uncertainty”). In this economy, changes in ϕ_{spread} have much smaller effects. Indeed, the standard deviation decreases from 0.64 only up to 0.63.

Second, we consider fiscal policy effects. In standard models, an increase in government spending crowds out consumption and hence the government spending multiplier on output, dY_t/dG_t , tends to be modest and below one. In our model, however, an increase in hours worked triggered by an increase in government spending raises agents’ confidence, which feeds back and raises the level of consumption and other economic activities. Because of this amplification effect, the government spending multiplier could be larger and above one.

In Figure 6, we plot the multiplier in our estimated model after a one-time, positive shock to government spending at $t = 1$.³⁸ The model predicts a multiplier that becomes larger than one after six quarters with a peak value at around 3.1. In contrast, in the RE counterfactual, where we set the entropy constraint to $\eta = 0$ and keep other parameters at their estimated values, the multiplier stays persistently below one.³⁹ We also plot the multipliers for various values of the

³⁸We assume that the government spending G_t in the resource constraint (2.12) is given by $G_t = g_t Y_t$, where g_t follows $\ln g_t = (1 - 0.95) \ln \bar{g} + 0.95 \ln g_{t-1} + \epsilon_{g,t}$.

³⁹We also computed a multiplier in the re-estimated RE model, where we set $\eta = 0$ and re-estimated the remaining parameters, and found that there the multiplier also stays below one.

entropy constraint in an increment of 0.02 from $\eta = 0$ to the estimated baseline value of $\eta = 1.7$ to further understand the role of confidence. Interestingly, the relationship between the size of the entropy constraint and the government spending multiplier is non-linear; the higher the entropy constraint, the larger effect an additional increase of the entropy constraint has on the size of the multiplier. An increase in government spending raises hours and hence confidence, which in turn feeds back to further raise hours and confidence. When ambiguity is larger, this feedback is stronger and hence the effect of an increase in the entropy constraint magnifies and has a stronger effect on the government spending multiplier.

It is important to emphasize that the large effects of government spending on output are not welfare increasing even though it arises due to a reduction in uncertainty. Indeed, since in this model learning arises at firm-level there are no information externalities that the government can correct. This is in contrast to models where learning occurs through observing the aggregate economy and it highlights the importance of modeling the underlying source of uncertainty for evaluating policies.

At a more general level, the comparisons of these counterfactual models in the monetary and fiscal policy experiments underscore the importance for policy analysis of modeling time-variation in uncertainty as an endogenous response that in turn further affects economic decisions.

5 Conclusion

In this paper we construct a tractable heterogeneous-firm business cycle model in which a representative household faces Knightian uncertainty about the firm level profitability. Firm's production serves as a signal about this hidden state and learning is more informative for larger production scales. The feedback loop between economic activity and confidence makes our model behave as a standard linear business cycle model with (i) countercyclical wedges in the equilibrium supply for labor, for risk-free as well as for risky assets, (ii) positive co-movement of aggregate variables in response to either supply or demand shocks, and (iii) strong internal propagation with amplified and hump-shaped dynamics. When the model is estimated using US macroeconomic and financial data, we find that (i) a financial shock emerges as a key source of business cycles, (ii) the empirical role of traditional frictions become smaller, and (iii) the aggregate activity becomes more responsive to monetary and fiscal policies. We conclude that endogenous idiosyncratic uncertainty is a quantitatively important mechanism for understanding business cycles.

References

- AN, S. AND F. SCHORFHEIDE (2007): “Bayesian Analysis of DSGE Models,” *Econometric Reviews*, 26, 113–172.
- ANGELETOS, G.-M., F. COLLARD, AND H. DELLAS (2014): “Quantifying Confidence,” *NBER Working Paper No. 20807*.
- ANGELETOS, G.-M. AND J. LA’O (2009): “Noisy business cycles,” in *NBER Macroeconomics Annual 2009, Volume 24*, 319–378.
- (2013): “Sentiments,” *Econometrica*, 81, 739–779.
- ARELLANO, C., Y. BAI, AND P. J. KEHOE (2012): “Financial Markets and Fluctuations in Uncertainty,” *Federal Reserve Bank of Minneapolis Staff Report*.
- BACHMANN, R. AND C. BAYER (2013): “‘Wait-and-See’ Business Cycles?” *Journal of Monetary Economics*, 60, 704–719.
- BARRO, R. J. AND R. G. KING (1984): “Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles,” *Quarterly Journal of Economics*, 99, 817–839.
- BASU, S. AND B. BUNDICK (2012): “Uncertainty Shocks in a Model of Effective Demand,” *NBER Working Paper No. 18420*.
- BERNANKE, B. S., M. GERTLER, AND S. GILCHRIST (1999): “The Financial Accelerator in a Quantitative Business Cycle Framework,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, Amsterdam: Elsevier, vol. 1, Part C, chap. 21, 1341–1393.
- BIANCHI, F., C. L. ILUT, AND M. SCHNEIDER (2014): “Uncertainty Shocks, Asset Supply and Pricing over the Business Cycle,” *NBER Working Paper No. 20081*.
- BLOOM, N. (2009): “The Impact of Uncertainty Shocks,” *Econometrica*, 77, 623–685.
- BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA-EKSTEN, AND S. TERRY (2014): “Really Uncertain Business Cycles,” *Working Paper*.
- BORN, B. AND J. PFEIFER (2014): “Policy Risk and the Business Cycle,” *Journal of Monetary Economics*, 68, 68–85.
- CAPLIN, A. AND J. LEAHY (1993): “Sectoral Shocks, Learning, and Aggregate Fluctuations,” *The Review of Economic Studies*, 60, 777–794.
- CHARI, V. V., P. J. KEHOE, AND E. R. MCGRATTAN (2007): “Business Cycle Accounting,” *Econometrica*, 75, 781–836.

- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113, 1–45.
- CHRISTIANO, L. J., M. S. EICHENBAUM, AND M. TRABANDT (2015): “Understanding the Great Recession,” *American Economic Journal: Macroeconomics*, 7, 110–167.
- CHRISTIANO, L. J., C. L. ILUT, R. MOTTO, AND M. ROSTAGNO (2010): “Monetary Policy and Stock Market Booms,” NBER Working Papers 16402.
- CHRISTIANO, L. J., R. MOTTO, AND M. ROSTAGNO (2014): “Risk Shocks,” *American Economic Review*, 104, 27–65.
- COCHRANE, J. H. (2011): “Presidential address: Discount rates,” *Journal of Finance*, 66, 1047–1108.
- CÚRDIA, V. AND M. WOODFORD (2010): “Credit Spreads and Monetary Policy,” *Journal of Monetary, Credit, and Banking*, 42, 3–35.
- DAVID, J. M., H. A. HOPENHAYN, AND V. VENKATESWARAN (2015): “Information, Misallocation, and Aggregate Productivity,” *Quarterly Journal of Economics*, forthcoming.
- EGGERTSSON, G. B. ET AL. (2003): “Zero bound on interest rates and optimal monetary policy,” *Brookings Papers on Economic Activity*, 2003, 139–233.
- EPSTEIN, L. G. AND M. SCHNEIDER (2003a): “IID: independently and indistinguishably distributed,” *Journal of Economic Theory*, 113, 32–50.
- (2003b): “Recursive Multiple-Priors,” *Journal of Economic Theory*, 113, 1–31.
- EUSEPI, S. AND B. PRESTON (2015): “Consumption Heterogeneity, Employment Dynamics and Macroeconomic Co-movement,” *Journal of Monetary Economics*, 71, 13–32.
- FAJGELBAUM, P., E. SCHAAL, AND M. TASCHEREAU-DUMOUCHEL (2016): “Uncertainty Traps,” *Quarterly Journal of Economics*, forthcoming.
- FERNÁNDEZ-VILLAYERDE, J., P. GUERRÓN-QUINTANA, K. KUESTER, AND J. F. RUBIO-RAMÍREZ (2015): “Fiscal Volatility Shocks and Economic Activity,” *American Economic Review*, forthcoming.
- FERNÁNDEZ-VILLAYERDE, J., P. GUERRÓN-QUINTANA, J. F. RUBIO-RAMÍREZ, AND M. URIBE (2011): “Risk Matters: The Real Effects of Volatility Shocks,” *American Economic Review*, 101, 2530–2561.
- FISHER, J. D. (2015): “On the Structural Interpretation of the Smets-Wouters Risk Premium Shock,” *Journal of Money, Credit and Banking*, 47, 511–516.

- GILBOA, I. AND D. SCHMEIDLER (1989): “Maxmin Expected Utility with Non-unique Prior,” *Journal of Mathematical Economics*, 18, 141–153.
- GILCHRIST, S. AND E. ZAKRAJŠEK (2012): “Credit Spreads and Business Cycle Fluctuations,” *American Economic Review*, 102, 1692–1720.
- GOURIO, F. (2014): “Financial distress and endogenous uncertainty,” Manuscript.
- HANSEN, G. D. (1985): “Indivisible Labor and the Business Cycle,” *Journal of Monetary Economics*, 16, 309–337.
- ILUT, C. AND M. SCHNEIDER (2014): “Ambiguous Business Cycles,” *American Economic Review*, 104, 2368–2399.
- JAIMOVICH, N. AND S. REBELO (2009): “Can News about the Future Drive the Business Cycle?” *American Economic Review*, 99, 1097–1118.
- KEHRIG, M. (2015): “The Cyclical Nature of Productivity Dispersion,” *Working Paper*.
- KOZLOWSKI, J., L. VELDKAMP, AND V. VENKATESWARAN (2015): “The tail that wags the economy: Belief-driven business cycles and persistent stagnation,” NBER Working Paper No. 21719.
- LORENZONI, G. (2009): “A Theory of Demand Shocks,” *American Economic Review*, 99, 2050–2084.
- MARINACCI, M. (1999): “Limit Laws for Non-Additive Probabilities, and Their Frequentist Interpretation,” *Journal of Economic Theory*, 84, 145–195.
- ORDOÑEZ, G. (2013): “The Asymmetric Effects of Financial Frictions,” *Journal of Political Economy*, 5, 844–895.
- ROGERSON, R. (1988): “Indivisible Labor, Lotteries and Equilibrium,” *Journal of Monetary Economics*, 21, 3–16.
- SAIJO, H. (2014): “The Uncertainty Multiplier and Business Cycles,” *Working Paper*.
- SENGA, T. (2015): “A New Look at Uncertainty Shock: Imperfect Information and Misallocation,” *Working Paper*.
- SHIMER, R. (2009): “Convergence in Macroeconomics: The Labor Wedge,” *American Economic Journal: Macroeconomics*, 1, 280–297.
- SIMS, C. A. (2002): “Solving Linear Rational Expectations Models,” *Computational Economics*, 20, 1–20.

SMETS, F. AND R. WOUTERS (2007): “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *The American Economic Review*, 97, 586–606.

STRAUB, L. AND R. ULBRICHT (2016): “Endogenous Uncertainty and Credit Crunches,” Manuscript.

TAYLOR, J. B. (2008): “Monetary Policy and the State of the Economy,” *Testimony before the Committee on Financial Services U.S. House of Representatives*.

VAN NIEUWERBURGH, S. AND L. VELDKAMP (2006): “Learning Asymmetries in Real Business Cycles,” *Journal of Monetary Economics*, 53, 753–772.

6 Appendix (For online publication)

6.1 A model of learning about firm-specific demand

In this section, we show that the baseline model with additive shock in the production function (2.5) can be reinterpreted as a model where firms learn about their demand from noisy signals.

There is a continuum of firms, indexed by $l \in [0, 1]$, which produce intermediate goods and sell them to a large representative “conglomerate”. The conglomerate, who holds shares of the intermediate firms, acts in a perfectly competitive manner and combine the intermediate goods to produce final goods. To ease exposition, we momentarily abstract from all aggregate shocks and labor-augmenting technological growth.

The conglomerate combines intermediate output $Y_{l,t}$ according to the following CES aggregator:

$$Y_t = \left[\int_0^1 z_{l,t}^{\frac{\sigma-1}{\sigma}} Y_{l,t}^{\frac{\sigma-1}{\sigma}} dl \right]^{\frac{\sigma}{\sigma-1}}, \quad \sigma > 1$$

where $z_{l,t}$ follows an AR(1) as in (2.4). In turn, intermediate goods $Y_{l,t}$ are produced according to

$$Y_{l,t} = K_{l,t-1}^\alpha H_{l,t}^{1-\alpha}.$$

We assume that all agents, including the conglomerate and households, cannot observe the realization of $z_{l,t}$. Instead, after the production of final and intermediate goods they observe noisy signals

$$s_{l,t} = Y_{l,t} z_{l,t} + \tilde{v}_{l,t}, \quad \tilde{v}_{l,t} \sim N(0, \sigma_v^2) \tag{6.1}$$

where $\tilde{v}_{l,t}$ is an observation error. Agents use all available information, including the path of signals $s_{l,t}$ and intermediate output $Y_{l,t}$, to form estimates about the realization of $z_{l,t}$.

In this context, it is natural to think $z_{l,t}$ as a “quality” of intermediate good l that is difficult to observe. Alternatively, since the final good could directly be used for consumption, $z_{l,t}$ can be

interpreted as an unobservable demand for a variety l . Crucially, (6.1) implies that the signal-to-noise ratio is increasing in the level of output $Y_{l,t}$. It is plausible that firms learn more about the quality or demand, $z_{l,t}$, of their goods when they produce and sell more. For example, when a restaurant serves more customers it generates more website reviews and hence people learn more about the quality of their meals.

Intermediate firms l choose price $P_{l,t}$ and inputs to maximize the shareholder value

$$E_0^* \sum_{t=0}^{\infty} M_0^t D_{l,t}.$$

$D_{l,t}$ is the dividend payout to the conglomerate given by

$$D_{l,t} = \frac{P_{l,t}}{P_t} Y_{l,t} - W_t H_{l,t} - I_{l,t},$$

where P_t is the price of the final good given by

$$P_t = \left[\int_0^1 P_{l,t}^{1-\sigma} dl \right]^{\frac{1}{1-\sigma}}.$$

The conglomerate, in turn, choose intermediate inputs $Y_{l,t}$ and shares $\theta_{l,t}$ to maximize the shareholder value

$$E_0^* \sum_{t=0}^{\infty} M_0^t D_t,$$

where D_t is the dividend payout to the households

$$D_t = Y_t + \int (D_{l,t} + P_{l,t}^e) \theta_{l,t-1} dl - \int P_{l,t}^e \theta_{l,t} dl.$$

The household side of the economy is the same as in the baseline model except that the households hold shares θ_t of conglomerates instead of shares $\theta_{l,t}$ of intermediate firms.

We now reintroduce aggregate shocks and utilization and describe the timing of the event at period t .

1. Stage 1 : Pre-production stage

- Agents observe the realization of aggregate shocks (A_t and Δ_t^k).
- Given forecasts about $z_{l,t}$ and its associated worst-case scenario, firms hire labor and choose price ($H_{l,t}$ and $P_{l,t}$). The household supply labor H_t and the labor market clears at the wage rate W_t .
- Firms produce intermediate output $Y_{l,t}$ and sells it to the conglomerates at price $P_{l,t}$.

2. Stage 2 : Post-production stage

- $z_{l,t}$ realize (but are unobservable) and production of the final goods Y_t takes place. Agents observe noisy signals $s_{l,t}$.
- Firms and conglomerates update estimates about $z_{l,t}$ and use it to form forecasts for production next period.
- Firms make investment $I_{l,t}$ and pay out dividends $D_{l,t}$ to the conglomerates. The conglomerates make asset purchase decisions $\theta_{l,t}$ and pay out dividends D_t to the households. Finally, households make consumption and asset purchase decisions (C_t , B_t , and θ_t).

In the perfect competition limit ($\sigma \rightarrow \infty$), this version of the model is observationally equivalent to the baseline model at the aggregate level. The introduction of the conglomerate is important for two reasons. First, it prohibits households from inferring $z_{l,t}$ from utility by directly consuming intermediate goods $Y_{l,t}$. Second, it generates countercyclical ex-post excess return on equity held by the household. This is because dividend payout by the conglomerate to the household (D_t) is based on the realized return on capital. Note that the dividend payout by the intermediate firms to the conglomerate ($D_{l,t}$) is not based on the realized return since the production and market clearing of $Y_{l,t}$ happens before the realization of $z_{l,t}$.

6.2 Recursive competitive equilibrium for the frictionless model

We collect exogenous aggregate state variables (such as aggregate TFP) in a vector X with a cumulative transition function $F(X'|X)$. The endogenous aggregate state is the distribution of firm-level variables. A firm's type is identified by the posterior mean estimate of productivity \tilde{z}_l , the posterior variance Σ_l , and its capital stock K_l . The worst-case TFP is not included because it is implied by the posterior mean and variance. We denote the cross-sectional distribution of firms' type by ξ_1 and ξ_2 . ξ_1 is a stage 1 distribution over $(\tilde{z}_l, \Sigma_l, K_l)$ and ξ_2 is a stage 2 distribution over $(\tilde{z}'_l, \Sigma'_l, K'_l)$. ξ'_1 , in turn, is a distribution over $(\tilde{z}'_l, \Sigma'_l, K'_l)$ at stage 1 in the next period.⁴⁰

First, consider the household's problem. The household's wealth can be summarized by a portfolio $\vec{\theta}_l$ which consists of share θ_l for each firm and the risk-less bond holdings B . We use V_1^h and V_2^h to denote the household's value function at stage 1 and stage 2, respectively. We use m to summarize the income available to the household at stage 2. The household's problem at stage 1 is

$$\begin{aligned}
 V_1^h(\vec{\theta}_l, B; \xi_1, X) &= \max_H \left\{ -\frac{H^{1+\phi}}{1+\phi} + E^*[V_2^h(\hat{m}; \hat{\xi}_2, X)] \right\} \\
 \text{s.t. } \hat{m} &= WH + RB + \int (\hat{D}_l + \hat{P}_l)\theta_l dl
 \end{aligned} \tag{6.2}$$

⁴⁰See also Senga (2015) for a recursive representation of an imperfect information heterogeneous-firm model with time-varying uncertainty.

where we momentarily use the *hat* symbol to indicate random variables that will be resolved at stage 2. The household's problem at stage 2 is

$$\begin{aligned}
V_2^h(m; \xi_2, X) &= \max_{C, \vec{\theta}_l''} \left\{ \ln C + \beta \int V_1^h(\vec{\theta}_l''; \xi_1', X') dF(X'|X) \right\} \\
\text{s.t. } & C + B' + \int P_l \theta_l' dl \leq m \\
& \xi_1' = \Gamma(\xi_2, X)
\end{aligned} \tag{6.3}$$

In problem (6.2), households choose labor supply based on the worst-case stage 2 value (recall that we use E^* to denote worst-case conditional expectations). The problem (6.3), in turn, describes the household's consumption and asset allocation problem given the realization of income and aggregate states. In particular, they take as given the law of motion of the next period's distribution $\xi_1' = \Gamma(\xi_2, X)$, which in equilibrium is consistent with the firm's policy function. Importantly, in contrast to the stage 2 problem, a law of motion that describes the evolution of ξ_2 from (ξ_1, X) is absent in the stage 1 problem. Indeed, if there is no ambiguity in the model, agents take as given the law of motion $\xi_2 = \Upsilon(\xi_1, X)$, which in equilibrium is consistent with the firm's policy function and the true data generating process of the firm-level TFP. Since agents are ambiguous about each firm's TFP process, they cannot settle on a single law of motion about the distribution of firms. Finally, the continuation value at stage 2 is governed by the transition density of aggregate exogenous states X .

Next, consider the firms' problem. We use v_1^f and v_2^f to denote the firm's value function at stage 1 and stage 2, respectively. Firm l 's problem at stage 1 is

$$\begin{aligned}
v_1^f(\tilde{z}_l, \Sigma_l, K_l; \xi_1, X) &= \max_{H_l} E^*[v_2^f(\hat{z}_l', \Sigma_l', K_l; \hat{\xi}_2, X)] \\
\text{s.t. } & \text{Updating rules (2.8) and (2.10)}
\end{aligned} \tag{6.4}$$

and firm l 's problem at stage 2 is

$$\begin{aligned}
v_2^f(\tilde{z}_l', \Sigma_l', K_l; \xi_2, X) &= \max_{I_l} \left\{ \lambda(Y_l - WH_l - I_l) + \beta \int v_1^f(\tilde{z}_l', \Sigma_l', K_l'; \xi_1', X') dF(X'|X) \right\} \\
\text{s.t. } & K_l' = (1 - \delta)K_l + I_l \\
& \xi_1' = \Gamma(\xi_2, X)
\end{aligned} \tag{6.5}$$

where we simplify the exposition by expressing a firm's value in terms of the marginal utility λ of the representative household. Similar to the household's problem, a firm's problem at stage 1 is to choose the labor demand so as to maximize the worst-case stage 2 value. Note that the posterior mean \tilde{z}_l' will be determined by the realization of output Y_l at stage 2 while the posterior variance Σ_l' is determined by Σ_l and the input level at stage 1. In problem (6.5), the firm then chooses investment taking as given the realization of output and the updated estimates of its productivity.

The recursive competitive equilibrium is therefore a collection of value functions, policy functions, and prices such that

1. Households and firms optimize; (6.2) – (6.5).
2. The labor market, goods market, and asset markets clear.
3. The law of motion $\xi'_1 = \Gamma(\xi_2, X)$ is induced by the firms' policy function $I_l(\tilde{z}'_l, \Sigma'_l, K_l; \xi_2, X)$.

6.3 Equilibrium conditions for the estimated model

As we describe below in Appendix 6.4, we express equilibrium conditions from the perspective of agents at both stage 1 and stage 2. At stage 1, we need not only equilibrium conditions for variable determined before production (such as utilization and hours), but also those for variables determined after production (such as consumption and investment). At stage 2, we treat variables determined before production as pre-determined. To do this, we index period t variables determined at stage 1 by $t - 1$ and period t variables determined at stage 2 by t . We then combine stage 1 and stage 2 equilibrium conditions by using the certainty equivalence property of linearized decision rules.

We scale the variables in order to introduce stationary:

$$c_t = \frac{C_t}{\gamma^t}, y_{l,t} = \frac{Y_{l,t}}{\gamma^t}, k_{l,t-1} = \frac{K_{l,t-1}}{\gamma^t}, i_{l,t} = \frac{I_{l,t}}{\gamma^t}, \tilde{y}_{l,t} = \frac{\tilde{Y}_{l,t-1}}{\gamma^t}, w_t = \frac{W_t}{\gamma^t}, \tilde{\lambda}_t = \gamma^t \lambda_t, \tilde{\mu}_{l,t} = \gamma^t \mu_{l,t},$$

where $\mu_{l,t}$ is the Lagrangian multiplier on the capital accumulation equation of firm l .

We first describe the stage 1 equilibrium conditions. An individual firm l 's problem is to choose $\{U_{l,t}, K_{l,t}, H_{l,t}, I_{l,t}\}$ to maximize

$$E_t^* \sum_{s=0}^{\infty} \beta^{t+s} \lambda_{t+s} [P_{t+s}^W Y_{l,t+s} - W_{t+s} H_{l,t+s} - I_{l,t+s} - a(U_{l,t+s}) K_{l,t+s-1}],$$

where P_t^W is the price of whole-sale goods produced by firms and λ_t , and its detrended counterpart $\tilde{\lambda}_t$, is the marginal utility of the representative household:

$$\tilde{\lambda}_t = \frac{\gamma}{c_t - bc_{t-1}} - \beta b E_t^* \frac{1}{\gamma c_{t+1} - bc_t}, \quad (6.6)$$

subject to the following three constraints. The first constraint is the production function:

$$y_{l,t} = A_t \tilde{y}_{l,t} \{E_{t-1}^* z_{l,t} + \nu_{l,t}\}, \quad (6.7)$$

where $\tilde{y}_{l,t}$ is the input,

$$\tilde{y}_{l,t} = (U_{l,t} k_{l,t-1})^\alpha H_{l,t}^{1-\alpha}. \quad (6.8)$$

The worst case TFP $E_t^* z_{l,t+1|t+1}$ is given by

$$E_t^* z_{l,t+1|t+1} = (1 - \rho_z)\bar{z} + \rho_z \tilde{z}_{l,t|t} - \eta_a \rho_z \sqrt{\Sigma_{l,t|t}}. \quad (6.9)$$

and the Kalman filter estimate $\tilde{z}_{l,t|t}$ evolves according to

$$\tilde{z}_{l,t|t} = \tilde{z}_{l,t|t-1} + Gain_{l,t}(y_{l,t}/(\tilde{y}_{l,t}A_t) - \tilde{z}_{l,t|t-1}), \quad (6.10)$$

where $Gain_{l,t}$ is the Kalman gain and is given by

$$Gain_{l,t} = \left[\frac{\Sigma_{l,t|t-1}}{\Sigma_{l,t|t-1} + \tilde{y}_{l,t}^{-1} \sigma_\nu^2} \right]. \quad (6.11)$$

The second constraint is the capital accumulation equation:

$$\gamma k_{l,t} = (1 - \delta)k_{l,t-1} + \left\{ 1 - \frac{\kappa}{2} \left(\frac{\gamma i_{l,t}}{i_{l,t-1}} - \gamma \right)^2 \right\} i_{l,t} \quad (6.12)$$

and the last constraint is the law of motion for posterior variance:

$$\Sigma_{l,t|t} = (1 - Gain_{l,t})\Sigma_{l,t|t-1}. \quad (6.13)$$

As described in the main text, firms take into account the impact of their input choice on worst-case probabilities.

The first-order necessary conditions for firms' input choices are as follows:

- FONC for $\Sigma_{l,t|t}$

$$\begin{aligned} \psi_{l,t} = & \beta E_t^* \left[\frac{1}{2} \tilde{\lambda}_{t+1} P_{t+1}^W A_{t+1} \eta_a \rho_z \Sigma_{l,t|t}^{-\frac{1}{2}} \tilde{y}_{l,t+1} \right. \\ & \left. + \psi_{l,t+1} \left\{ \frac{\sigma_\nu^2 \rho_z^2}{\tilde{y}_{l,t+1}^2 (\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) + \sigma_\nu^2} - \frac{\sigma_\nu^2 \rho_z^2 (\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) \tilde{y}_{l,t+1}^2}{\{f_{l,t+1}^2 (\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) + \sigma_\nu^2\}^2} \right\} \right], \end{aligned} \quad (6.14)$$

where $\psi_{l,t}$ is the Lagrangian multiplier for the law of motion of posterior variance.

- FONC for $U_{l,t}$

$$\begin{aligned} & \tilde{\lambda}_t P_t^W \alpha \frac{y_{l,t}}{U_{l,t}} + \psi_{l,t} \frac{2\alpha \sigma_\nu^2 (\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) \tilde{y}_{l,t}^2}{\{\tilde{y}_{l,t}^2 (\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) + \sigma_\nu^2\}^2 U_{l,t}} \\ & = \tilde{\lambda}_t \{\chi_1 \chi_2 U_{l,t} + \chi_2 (1 - \chi_1)\} k_{l,t-1} \end{aligned} \quad (6.15)$$

- FONC for $H_{l,t}$

$$\tilde{\lambda}_t (1 - \alpha) P_t^W \frac{y_{l,t}}{H_{l,t}} + \psi_{l,t} \frac{2(1 - \alpha) \sigma_\nu^2 (\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) \tilde{y}_{l,t}^2}{\{\tilde{y}_{l,t}^2 (\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) + \sigma_\nu^2\}^2 H_{l,t}} = \tilde{\lambda}_t \tilde{w}_t, \quad (6.16)$$

where \tilde{w}_t is the real wage: $\tilde{w}_t \equiv w_t/P_t$.

- FONC for $k_{l,t}$

$$\gamma \tilde{\lambda}_t = \beta E_t^* (1 - \Delta_t^k) \tilde{\lambda}_{t+1} \frac{R_{t+1}^k}{\pi_{t+1}},$$

where the return on capital is defined as

$$R_t^k = \left[P_t^W \alpha \frac{y_{l,t}}{k_{l,t-1}} + q_{l,t}(1 - \delta) - a(U_{l,t}) \right. \\ \left. + \psi_{l,t}^k \frac{2\alpha\sigma_v^2(\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) \tilde{y}_{l,t}^2}{\{\tilde{y}_{l,t}^2(\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) + \sigma_v^2\}^2 k_{l,t-1}} \right] \times \frac{\pi_t}{q_{l,t-1}}, \quad (6.17)$$

where

$$q_{l,t} = \tilde{\mu}_{l,t} / \tilde{\lambda}_t, \quad (6.18)$$

$$\psi_{l,t}^k = \psi_{l,t} / \tilde{\lambda}_t. \quad (6.19)$$

- FONC for $i_{l,t}$

$$\gamma \tilde{\lambda}_t = \gamma \tilde{\mu}_{l,t} \left[1 - \frac{\kappa}{2} \left(\frac{\gamma i_{l,t}}{i_{l,t-1}} - \gamma \right)^2 - \kappa \left(\frac{\gamma i_{l,t}}{i_{l,t-1}} - \gamma \right) \frac{\gamma i_{l,t}}{i_{l,t-1}} \right] \\ + \beta E_t^* \left[\tilde{\mu}_{l,t+1} \kappa \left(\frac{\gamma i_{l,t+1}}{i_{l,t}} - \gamma \right) \left(\frac{\gamma i_{l,t+1}}{i_{l,t}} \right)^2 \right] \quad (6.20)$$

We eliminate l -subscripts to denote cross-sectional means (e.g., $y_t \equiv \int_0^1 y_{l,t} dl$).

Firms sell their wholesale goods to monopolistically competitive retailers. Conditions associated with Calvo sticky prices are

$$P_t^n = \tilde{\lambda}_t P_t^W y_t + \xi_p \beta E_t^* \left(\frac{\pi_{t+1}}{\bar{\pi}} \right)^{\theta_p} P_{t+1}^n \quad (6.21)$$

$$P_t^d = \tilde{\lambda}_t y_t + \xi_p \beta E_t^* \left(\frac{\pi_{t+1}}{\bar{\pi}} \right)^{\theta_p - 1} P_{t+1}^d \quad (6.22)$$

$$p_t^* = \left(\frac{\theta_p}{\theta_p - 1} \right) \frac{P_t^n}{P_t^d} \quad (6.23)$$

$$1 = (1 - \xi_p)(p_t^*)^{1 - \theta_p} + \xi_p \left(\frac{\bar{\pi}}{\pi_t} \right)^{1 - \theta_p} \quad (6.24)$$

$$y_t^* = \tilde{p}_t^{-\theta_p} y_t \quad (6.25)$$

$$\tilde{p}_t = (1 - \xi_p)(p_t^*)^{-\theta_p} + \xi_p \left(\frac{\bar{\pi}}{\pi_t} \right)^{-\theta_p} \quad (6.26)$$

Conditions associated with Calvo sticky wages are

$$v_t^1 = v_t^2 \quad (6.27)$$

$$v_t^1 = (w_t^*)^{1-\theta_w} \tilde{\lambda}_t H_t \tilde{w}_t + \xi_w \beta E_t^* \left(\frac{\pi_{t+1}^w w_{t+1}^*}{\bar{\pi} w_t^*} \right)^{\theta_w-1} v_{t+1}^1 \quad (6.28)$$

$$v_t^2 = \frac{\theta_w}{\theta_w - 1} (w_t^*)^{-\theta_w(1+\eta)} H_t^{1+\eta} + \xi_w \beta E_t^* \left(\frac{\pi_{t+1}^w w_{t+1}^*}{\bar{\pi} w_t^*} \right)^{\theta_w(1+\eta)} v_{t+1}^2 \quad (6.29)$$

$$1 = (1 - \xi_w) (w_t^*)^{1-\theta_w} + \xi_w E_t^* \left(\frac{\bar{\pi}}{\pi_t^w} \right)^{1-\theta_w} \quad (6.30)$$

$$\pi_t^w = \pi_t \tilde{w}_t / \tilde{w}_{t-1} \quad (6.31)$$

Households' Euler equation:

$$\gamma \tilde{\lambda}_t = \beta E_t^* \tilde{\lambda}_{t+1} R_t / \pi_{t+1} \quad (6.32)$$

Monetary policy rule:

$$\frac{R_t}{\bar{R}} = \left(\frac{R_{t-1}}{\bar{R}} \right)^{\rho_R} \left[\left(\frac{\pi_t}{\bar{\pi}} \right)^{\phi_\pi} \left(\frac{y_t}{\bar{y}} \right)^{\phi_Y} \right]^{1-\rho_R} \epsilon_{R,t} \quad (6.33)$$

Resource constraint:

$$c_t + i_t = (1 - \bar{g}) y_t \quad (6.34)$$

The 30 endogenous variables we solve are:

$$k_t, y_t, i_t, c_t, H_t, U_t, \tilde{y}_t, \tilde{\lambda}_t, \tilde{\mu}_t, \psi_t, R_t, R_t^k, q_t, \psi_t^k, E_t^* z_{t+1}, \tilde{z}_{t|t}, \Sigma_{t|t}, Gain_t, \\ P_t^W, P_t^n, P_t^d, p_t^*, \pi_t, y_t^*, \tilde{p}_t, v_t^1, v_t^2, \tilde{w}_t, w_t^*, \pi_t^w$$

We have listed 30 conditions above, from (6.6) to (6.34). Of the above 30 endogenous variables, those that are determined at stage 1 are:

$$H_t, U_t, \tilde{y}_t, v_t^1, v_t^2, \tilde{w}_t, w_t^*, \pi_t^w$$

We now describe the state 2 equilibrium conditions. To avoid repetitions, we only list conditions that are different from the state 1 conditions.

- (6.7):

$$y_{l,t} = \tilde{y}_{l,t-1} A_t \{ z_{l,t} + \nu_{l,t} \},$$

- (6.8):

$$\tilde{y}_{l,t} = (U_{l,t} k_{l,t})^\alpha H_{l,t}^{1-\alpha}$$

- (6.10):

$$\tilde{z}_{l,t|t} = \tilde{z}_{l,t|t-1} + Gain_{l,t}(y_{l,t}/A_t - \tilde{z}_{l,t|t-1})$$

- (6.11):

$$Gain_{l,t} = \left[\frac{\Sigma_{l,t|t-1}}{\Sigma_{l,t|t-1} + \tilde{y}_{l,t-1}\sigma_\nu^2} \right]$$

- (6.13):

$$\Sigma_{l,t|t} = (1 - Gain_{l,t})\Sigma_{l,t|t-1}$$

- (6.14):

$$\begin{aligned} \psi_{l,t} = & \beta E_t^* \left[\frac{1}{2} \tilde{\lambda}_{t+1} P_{t+1}^W A_{t+1} \eta_a \rho_z \Sigma_{l,t|t}^{-\frac{1}{2}} \tilde{y}_{l,t} \right. \\ & \left. + \psi_{l,t+1} \left\{ \frac{\sigma_\nu^2 \rho_z^2}{\tilde{y}_{l,t}^2 (\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) + \sigma_\nu^2} - \frac{\sigma_\nu^2 \rho_z^2 (\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) \tilde{y}_{l,t}^2}{\{\tilde{y}_{l,t}^2 (\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) + \sigma_\nu^2\}^2} \right\} \right] \end{aligned}$$

- (6.15):

$$\begin{aligned} & E_t^* \left[\tilde{\lambda}_{t+1} P_{t+1}^W \alpha \frac{y_{l,t+1}}{U_{l,t}} + \psi_{l,t+1} \frac{2\alpha \sigma_\nu^2 (\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) \tilde{y}_{l,t}^2}{\{\tilde{y}_{l,t}^2 (\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) + \sigma_\nu^2\}^2 U_{l,t}} \right] \\ & = E_t^* \tilde{\lambda}_{t+1} \{\chi_1 \chi_2 U_{l,t} + \chi_2 (1 - \chi_1)\} k_{l,t} \end{aligned}$$

- (6.16):

$$E_t^* \left[\tilde{\lambda}_{t+1} (1 - \alpha) P_{t+1}^W \frac{y_{l,t+1}}{H_{l,t}} + \psi_{l,t+1} \frac{2(1 - \alpha) \sigma_\nu^2 (\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) \tilde{y}_{l,t}^2}{\{\tilde{y}_{l,t}^2 (\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) + \sigma_\nu^2\}^2 H_{l,t}} \right] = E_t^* \tilde{\lambda}_{t+1} \tilde{w}_t$$

- (6.17):

$$\begin{aligned} R_t^k = & \left[P_t^W \alpha \frac{y_{l,t}}{k_{l,t-1}} + q_{l,t} (1 - \delta) - a(U_{l,t-1}) \right. \\ & \left. + \psi_{l,t}^k \frac{2\alpha \sigma_\nu^2 (\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) \tilde{y}_{l,t-1}^2}{\{\tilde{y}_{l,t-1}^2 (\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) + \sigma_\nu^2\}^2 k_{l,t-1}} \right] \times \frac{\pi_t}{q_{l,t-1}} \end{aligned}$$

- (6.28):

$$v_t^1 = (w_t^*)^{1-\theta_w} E_t^* \tilde{\lambda}_{t+1} H_t \tilde{w}_t + \xi_w \beta E_t^* \left(\frac{\pi_{t+1}^w w_{t+1}^*}{\bar{\pi} w_t^*} \right)^{\theta_w - 1} v_{t+1}^1$$

- (6.31):

$$\pi_t^w = E_t^* \pi_{t+1} \tilde{w}_t / \tilde{w}_{t-1}$$

6.4 Solution procedure

Here we describe the general solution procedure of the model. The procedure follows the method used in the example in Section 2.5. First, we derive the law of motion assuming that the model is a rational expectations model where the worst case expectations are on average correct. Second, we take the equilibrium law of motion formed under ambiguity and then evaluate the dynamics under the econometrician's data generating process. We provide a step-by-step description of the procedure:

1. Find the worst-case steady state.

We first compute the steady state of the filtering problem (2.8), (2.9), (2.10), and (2.15), under the worst-case mean (2.19) to find the firm-level TFP at the worst-case steady state, \bar{z}^0 . We then solve the steady state for other equilibrium conditions evaluated at \bar{z}^0 .

2. Log-linearize the model around the worst-case steady state.

We can solve for the dynamics using standard tools for linear rational expectation models. We base our discussion based on the method proposed by Sims (2002).

We first need to deal with the issue that idiosyncratic shocks realize at the beginning of stage 2. Handling this issue correctly is important, since variables chosen at stage 1, such as input choice, should be based on the worst-case TFP, while variables chosen at stage 2, such as consumption and investment, would be based on the realized TFP (but also on the worst-case future TFP). To do this, we exploit the certainty equivalence property of linear decision rules. We first solve for decision rules *as if* both aggregate and idiosyncratic shocks realize at the beginning of the period. We call them “pre-production decision rules”. We then solve for decision rules *as if* (i) both aggregate and idiosyncratic shocks realize at the beginning of the period and (ii) stage 1 variables are pre-determined. We call them “post-production decision rules”. Finally, when we characterize the dynamics from the perspective of the econometrician, we combine the pre-production and post-production decision rules and obtain an equilibrium law of motion.

To obtain pre-production decision rules, we collect the linearized equilibrium conditions, which include firm-level conditions, into the canonical form:

$$\mathbf{\Gamma}_0^{pre} \hat{\mathbf{y}}_t^{pre,0} = \mathbf{\Gamma}_1^{pre} \hat{\mathbf{y}}_{t-1}^{pre,0} + \mathbf{\Psi}^{pre} \omega_t + \mathbf{\Upsilon}^{pre} \eta_t^{pre},$$

where $\hat{\mathbf{y}}_t^{pre,0}$ is a column vector of size k that contains all variables and the conditional expectations. $\hat{\mathbf{y}}_t^{pre,0} = \mathbf{y}_t^{pre} - \bar{\mathbf{y}}^0$ denotes deviations from the worst-case steady state and η_t are expectation errors, which we define as $\eta_t^{pre} = \hat{\mathbf{y}}_t^{pre,0} - E_{t-1}^* \hat{\mathbf{y}}_t^{pre,0}$ such that $E_{t-1}^* \eta_t^{pre} = 0$. We define $\omega_t = [e_{l,t} \quad e_t]'$, where $e_{l,t} = [\epsilon_{z,l,t} \quad u_{l,t} \quad \nu_{l,t}]'$ is a vector of idiosyncratic shocks and

e_t is a vector of aggregate shocks of size n . For example, for the baseline model introduced in Section 2, e_t is a 2×1 vector of aggregate TFP and financial shocks.

The vector $\hat{\mathbf{y}}_t^{pre,0}$ contains firm-level variables such as firm l 's labor input, $H_{l,t}$. In contrast to other linear heterogeneous-agent models with imperfect information such as Lorenzoni (2009), all agents share the same information set. Thus, to derive the aggregate law of motion, we simply aggregate over firm l 's linearized conditions and replace firm-specific variables with their cross-sectional means (e.g., we replace $H_{l,t}$ with $H_t \equiv \int_0^1 H_{l,t} dl$) and set $e_{l,t} = \mathbf{0}$, which uses the law of large numbers for idiosyncratic shocks.

We order variables in $\hat{\mathbf{y}}_t^{pre,0}$ as

$$\hat{\mathbf{y}}_t^{pre,0} = \begin{bmatrix} \hat{\mathbf{y}}_{1,t}^{pre,0} \\ \hat{\mathbf{y}}_{2,t}^{pre,0} \\ \hat{\mathbf{s}}_t^{pre,0} \end{bmatrix},$$

where $\hat{\mathbf{y}}_{1,t}^{pre,0}$ is a column vector of size k_1 of variables determined at stage 1, $\hat{\mathbf{y}}_{2,t}^{pre,0}$ is a column vector of size k_2 of variables determined at stage 2, and $\hat{\mathbf{s}}_t^{pre,0} = [\hat{s}_{1,t}^{pre,0} \quad \hat{s}_{2,t}^{pre,0}]'$, where $s_{1,t} = \bar{z} - E_{t-1}^* z_t$ and $s_{2,t} = \bar{z} - \tilde{z}_{t|t}$.

The resulting solution of pre-production decision rules is obtained applying the method developed by Sims (2002):

$$\hat{\mathbf{y}}_t^{pre,0} = \mathbf{T}^{pre} \hat{\mathbf{y}}_{t-1}^{pre,0} + \mathbf{R}^{pre} [\mathbf{0}_{3 \times 1} \quad e_t]', \quad (6.35)$$

where \mathbf{T}^{pre} and \mathbf{R}^{pre} are $k \times k$ and $k \times (n + 3)$ matrices, respectively.

The solution of post-production decision rules can be obtained in a similar way by first collecting the equilibrium conditions into the canonical form

$$\mathbf{\Gamma}_0^{post} \hat{\mathbf{y}}_t^{post,0} = \mathbf{\Gamma}_1^{post} \hat{\mathbf{y}}_{t-1}^{post,0} + \mathbf{\Psi}^{post} \omega_t + \mathbf{\Upsilon}^{post} \eta_t^{post},$$

and is given by

$$\hat{\mathbf{y}}_t^{post,0} = \mathbf{T}^{post} \hat{\mathbf{y}}_{t-1}^{post,0} + \mathbf{R}^{post} [\mathbf{0}_{3 \times 1} \quad e_t]', \quad (6.36)$$

where

$$\hat{\mathbf{y}}_t^{post,0} = \begin{bmatrix} \hat{\mathbf{y}}_{1,t}^{post,0} \\ \hat{\mathbf{y}}_{2,t}^{post,0} \\ \hat{\mathbf{s}}_t^{post,0} \end{bmatrix},$$

and \mathbf{T}^{post} and \mathbf{R}^{post} are $k \times k$ and $k \times (n + 3)$ matrices, respectively.

3. Characterize the dynamics from the econometrician's perspective.

The above law of motion was based on the worst-case probabilities. We need to derive the equilibrium dynamics under the true DGP, where the cross-sectional mean of firm-level TFP

is \bar{z} . We are interested in two objects: the zero-risk steady state and the dynamics around that zero-risk steady state.

(a) Find the zero-risk steady state.

This the fixed point $\bar{\mathbf{y}}$ where the decision rules (6.35) and (6.36) are evaluated at the realized cross-sectional mean of firm-level TFP \bar{z} :

$$\begin{aligned}\bar{\mathbf{y}}^{pre} - \bar{\mathbf{y}}^0 &= \mathbf{T}^{pre}(\bar{\mathbf{y}} - \bar{\mathbf{y}}^0), \\ \bar{\mathbf{y}}^{post} - \bar{\mathbf{y}}^0 &= \mathbf{T}^{post}(\bar{\mathbf{y}} - \bar{\mathbf{y}}^0) + \mathbf{R}^{post}[\bar{\mathbf{s}} \quad \mathbf{0}_{(n+1) \times 1}]',\end{aligned}\tag{6.37}$$

where

$$\bar{\mathbf{y}} = \begin{bmatrix} \bar{\mathbf{y}}_1^{pre} \\ \bar{\mathbf{y}}_2^{post} \\ \bar{\mathbf{s}}^{post} \end{bmatrix}.$$

Note that we do not feed in the realized firm-level TFP to the pre-production decision rules since idiosyncratic shocks realize at the beginning of stage 2.

We obtain $\bar{\mathbf{s}}$ from

$$\bar{\mathbf{s}} = [\mathbf{T}_{3,1}^{post} \quad \mathbf{T}_{3,2}^{post} \quad \mathbf{T}_{3,3}^{post}](\bar{\mathbf{y}} - \bar{\mathbf{y}}^0) + \bar{\mathbf{s}}^0,$$

where $\mathbf{T}_{3,1}^{post}$, $\mathbf{T}_{3,2}^{post}$, and $\mathbf{T}_{3,3}^{post}$ are $2 \times k_1$, $2 \times k_2$, and 2×2 submatrices of \mathbf{T}^{post} , respectively:

$$\mathbf{T}^{post} = \begin{bmatrix} \mathbf{T}_{1,1}^{post} & \mathbf{T}_{1,2}^{post} & \mathbf{T}_{1,3}^{post} \\ \mathbf{T}_{2,1}^{post} & \mathbf{T}_{2,2}^{post} & \mathbf{T}_{2,3}^{post} \\ \mathbf{T}_{3,1}^{post} & \mathbf{T}_{3,2}^{post} & \mathbf{T}_{3,3}^{post} \end{bmatrix},$$

where $\mathbf{T}_{1,1}^{post}$, $\mathbf{T}_{1,2}^{post}$, $\mathbf{T}_{1,3}^{post}$, $\mathbf{T}_{2,1}^{post}$, $\mathbf{T}_{2,2}^{post}$, and $\mathbf{T}_{2,3}^{post}$ are $k_1 \times k_1$, $k_1 \times k_2$, $k_1 \times 2$, $k_2 \times k_1$, $k_2 \times k_2$, and $k_2 \times 2$ matrices, respectively.

(b) Dynamics around the zero-risk steady state.

Denoting $\hat{\mathbf{y}}_t \equiv \mathbf{y}_t - \bar{\mathbf{y}}$ the deviations from the zero-risk steady state, we combine the decision rules (6.35) and (6.36) evaluated at the true DGP and the equations for the zero-risk steady state (6.37) to characterize the equilibrium law of motion:

$$\hat{\mathbf{y}}_t^{pre} = \mathbf{T}^{pre}\hat{\mathbf{y}}_{t-1} + \mathbf{R}^{pre}[\mathbf{0}_{3 \times 1} \quad e_t]',\tag{6.38}$$

$$\hat{\mathbf{y}}_t^{post} = \mathbf{T}^{post}[\hat{\mathbf{y}}_{1,t}^{pre} \quad \hat{\mathbf{y}}_{2,t-1} \quad \hat{\mathbf{s}}_{t-1}]' + \mathbf{R}^{post}[\hat{\mathbf{s}}_t \quad 0 \quad e_t]',\tag{6.39}$$

$$\hat{\mathbf{s}}_t = [\mathbf{T}_{3,1}^{post} \quad \mathbf{T}_{3,2}^{post} \quad \mathbf{T}_{3,3}^{post}][\hat{\mathbf{y}}_{1,t}^{pre} \quad \hat{\mathbf{y}}_{2,t-1} \quad \hat{\mathbf{s}}_{t-1}]' + \mathbf{R}_{3,3}^{post}[\mathbf{0}_{3 \times 1} \quad e_t]',\tag{6.40}$$

and

$$\hat{\mathbf{y}}_t = \begin{bmatrix} \hat{\mathbf{y}}_{1,t}^{pre} \\ \hat{\mathbf{y}}_{2,t}^{post} \\ \hat{\mathbf{s}}_t^{post} \end{bmatrix}.\tag{6.41}$$

$\mathbf{R}_{3,3}^{post}$ is a $2 \times n$ submatrix of \mathbf{R}^{post} :

$$\mathbf{R}^{post} = \begin{bmatrix} \mathbf{R}_{1,1}^{post} & \mathbf{R}_{1,2}^{post} & \mathbf{R}_{1,3}^{post} \\ \mathbf{R}_{2,1}^{post} & \mathbf{R}_{2,2}^{post} & \mathbf{R}_{2,3}^{post} \\ \mathbf{R}_{3,1}^{post} & \mathbf{R}_{3,2}^{post} & \mathbf{R}_{3,3}^{post} \end{bmatrix},$$

where $\mathbf{R}_{1,1}^{post}$, $\mathbf{R}_{1,2}^{post}$, $\mathbf{R}_{1,3}^{post}$, $\mathbf{R}_{2,1}^{post}$, $\mathbf{R}_{2,2}^{post}$, $\mathbf{R}_{2,3}^{post}$, $\mathbf{R}_{3,1}^{post}$, and $\mathbf{R}_{3,2}^{post}$ are $k_1 \times 2$, $k_1 \times 1$, $k_1 \times n$, $k_2 \times 2$, $k_2 \times 1$, $k_2 \times n$, 2×2 , and 2×1 matrices, respectively.

We combine equations (6.38), (6.39), (6.40), and (6.41) to obtain the equilibrium law of motion. To do so, we first define submatrices of \mathbf{T}^{pre} and \mathbf{R}^{pre} :

$$\mathbf{T}^{pre} = \begin{bmatrix} \mathbf{T}_1^{pre} \\ \mathbf{T}_2^{pre} \\ \mathbf{T}_3^{pre} \end{bmatrix},$$

where \mathbf{T}_1^{pre} , \mathbf{T}_2^{pre} , and \mathbf{T}_3^{pre} are $k_1 \times k$, $k_2 \times k$, and $2 \times k$ matrices, respectively, and

$$\mathbf{R}^{pre} = \begin{bmatrix} \mathbf{R}_{1,1}^{pre} & \mathbf{R}_{1,2}^{pre} \\ \mathbf{R}_{2,1}^{pre} & \mathbf{R}_{2,2}^{pre} \\ \mathbf{R}_{3,1}^{pre} & \mathbf{R}_{3,2}^{pre} \end{bmatrix},$$

where $\mathbf{R}_{1,1}^{pre}$, $\mathbf{R}_{1,2}^{pre}$, $\mathbf{R}_{2,1}^{pre}$, $\mathbf{R}_{2,2}^{pre}$, $\mathbf{R}_{3,1}^{pre}$, and $\mathbf{R}_{3,2}^{pre}$ are $k_1 \times 3$, $k_1 \times n$, $k_2 \times 3$, $k_2 \times n$, 2×3 , and $2 \times n$ matrices, respectively.

We then define matrices \mathbf{T} and \mathbf{R} . A $k \times k$ matrix \mathbf{T} is given by

$$\mathbf{T} = \begin{bmatrix} \mathbf{T}_1^{pre} \\ \mathbf{T}_2 \\ \mathbf{T}_3 \end{bmatrix},$$

where \mathbf{T}_2 and \mathbf{T}_3 are $k_2 \times k_2$ and 2×2 matrices, respectively, given by

$$\begin{aligned} \mathbf{T}_2 &= [\mathbf{Q}_{2,1} \quad \mathbf{Q}_{2,2} + \mathbf{T}_{2,2}^{post} + \mathbf{R}_{2,1}^{post} \mathbf{T}_{3,2}^{post} \quad \mathbf{Q}_{2,3} + \mathbf{T}_{2,3}^{post} + \mathbf{R}_{2,1}^{post} \mathbf{T}_{3,3}^{post}], \\ \mathbf{T}_3 &= [\mathbf{Q}_{3,1} \quad \mathbf{Q}_{3,2} + \mathbf{T}_{3,2}^{post} + \mathbf{R}_{3,1}^{post} \mathbf{T}_{3,2}^{post} \quad \mathbf{Q}_{3,3} + \mathbf{T}_{3,3}^{post} + \mathbf{R}_{3,1}^{post} \mathbf{T}_{3,3}^{post}], \end{aligned}$$

and $\mathbf{Q}_{2,1}$, $\mathbf{Q}_{2,2}$, and $\mathbf{Q}_{2,3}$ are $k_2 \times k_1$, $k_2 \times k_2$, and $k_2 \times 2$ submatrices of \mathbf{Q}_2 , where $\mathbf{Q}_2 \equiv (\mathbf{T}_{2,1}^{post} + \mathbf{R}_{2,1}^{post} \mathbf{T}_{3,1}^{post}) \mathbf{T}_1^{pre}$, so that $\mathbf{Q}_2 = [\mathbf{Q}_{2,1} \quad \mathbf{Q}_{2,2} \quad \mathbf{Q}_{2,3}]$. Similarly, $\mathbf{Q}_{3,1}$, $\mathbf{Q}_{3,2}$, and $\mathbf{Q}_{3,3}$ are $k_3 \times k_1$, $k_3 \times k_2$, and $k_3 \times 2$ submatrices of \mathbf{Q}_3 , where $\mathbf{Q}_3 \equiv (\mathbf{T}_{3,1}^{post} + \mathbf{R}_{3,1}^{post} \mathbf{T}_{3,1}^{post}) \mathbf{T}_1^{pre}$, so that $\mathbf{Q}_3 = [\mathbf{Q}_{3,1} \quad \mathbf{Q}_{3,2} \quad \mathbf{Q}_{3,3}]$.

A $k \times n$ matrix \mathbf{R} is given by

$$\mathbf{R} = \begin{bmatrix} \mathbf{R}_{1,2}^{pre} \\ \mathbf{R}_2 \\ \mathbf{R}_3 \end{bmatrix},$$

where

$$\begin{aligned} \mathbf{R}_2 &= \mathbf{T}_{2,1}^{post} \mathbf{R}_{1,2}^{pre} + \mathbf{R}_{2,1}^{post} (\mathbf{T}_{3,1}^{post} \mathbf{R}_{1,2}^{pre} + \mathbf{R}_{3,3}^{post}) + \mathbf{R}_{2,3}^{post}, \\ \mathbf{R}_3 &= \mathbf{T}_{3,1}^{post} \mathbf{R}_{1,2}^{pre} + \mathbf{R}_{3,1}^{post} (\mathbf{T}_{3,1}^{post} \mathbf{R}_{1,2}^{pre} + \mathbf{R}_{3,3}^{post}) + \mathbf{R}_{3,3}^{post}. \end{aligned}$$

The equilibrium law of motion is then given by

$$\hat{\mathbf{y}}_t = \mathbf{T} \hat{\mathbf{y}}_{t-1} + \mathbf{R} e_t.$$

6.5 Data sources

We use the following data:

1. Real GDP in chained dollars, BEA, NIPA table 1.1.6, line 1.
2. GDP, BEA, NIPA table 1.1.5, line 1.
3. Personal consumption expenditures on nondurables, BEA, NIPA table 1.1.5, line 5.
4. Personal consumption expenditures on services, BEA, NIPA table 1.1.5, line 6.
5. Gross private domestic fixed investment (nonresidential and residential), BEA, NIPA table 1.1.5, line 8.
6. Personal consumption expenditures on durable goods, BEA, NIPA table 1.1.5, line 4.
7. Nonfarm business hours worked, BLS PRS85006033.
8. Civilian noninstitutional population (16 years and over), BLS LNU00000000.
9. Effective federal funds rate, Board of Governors of the Federal Reserve System.
10. Moody's Seasoned Baa Corporate Bond Yield Relative to Yield on 10-Year Treasury Constant Maturity (Baa spread), downloaded from Federal Reserve Economic Data, Federal Reserve Bank of St. Louis.

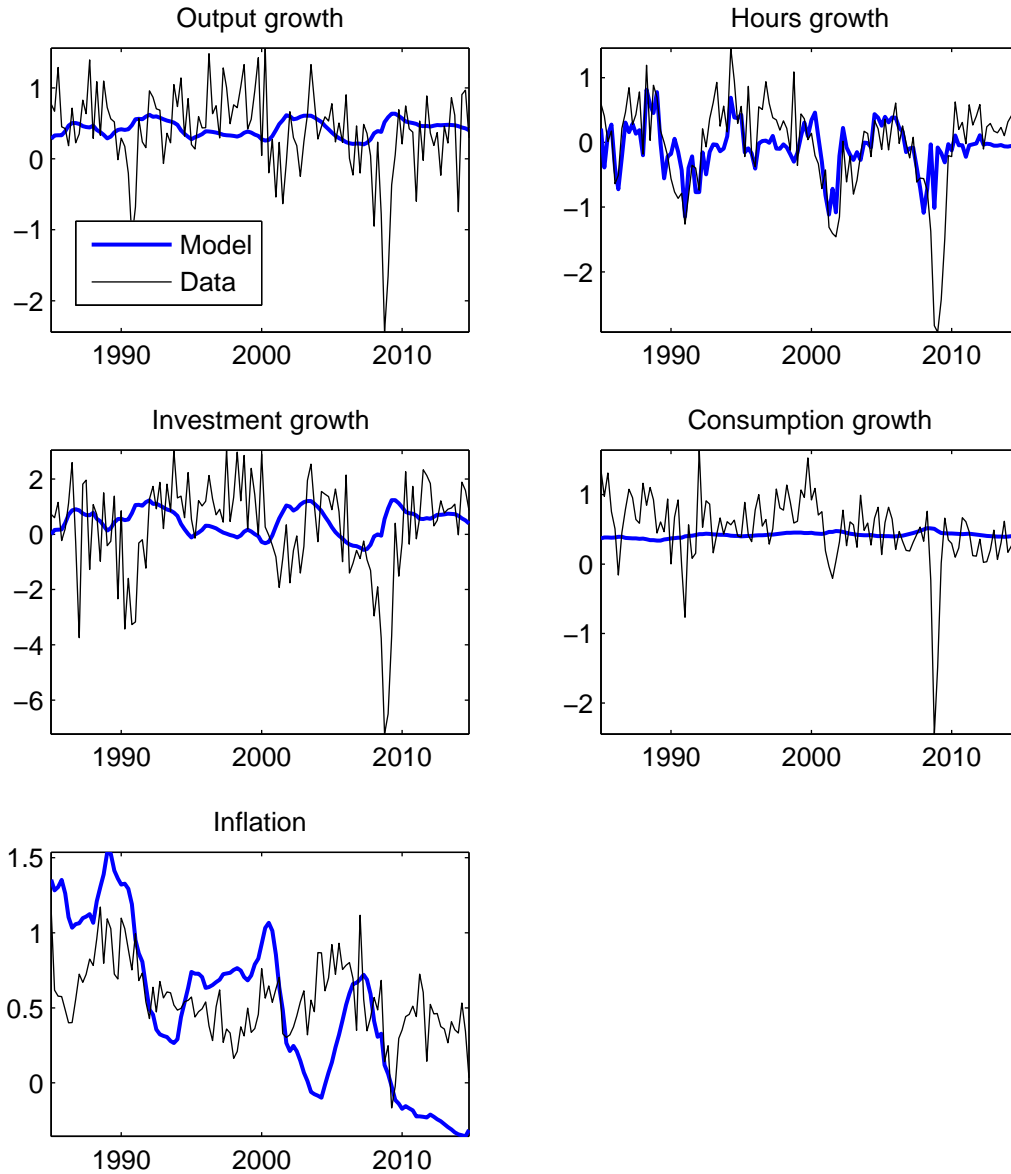
We then conduct the following transformations of the above data:

11. Real per capita GDP: (1)/(8)

12. GDP deflator: $(2)/(1)$
13. Real per capita consumption: $[(3)+(4)]/[(8)\times(12)]$
14. Real per capita investment: $[(5)+(6)]/[(8)\times(12)]$
15. Per capita hours: $(7)/(8)$

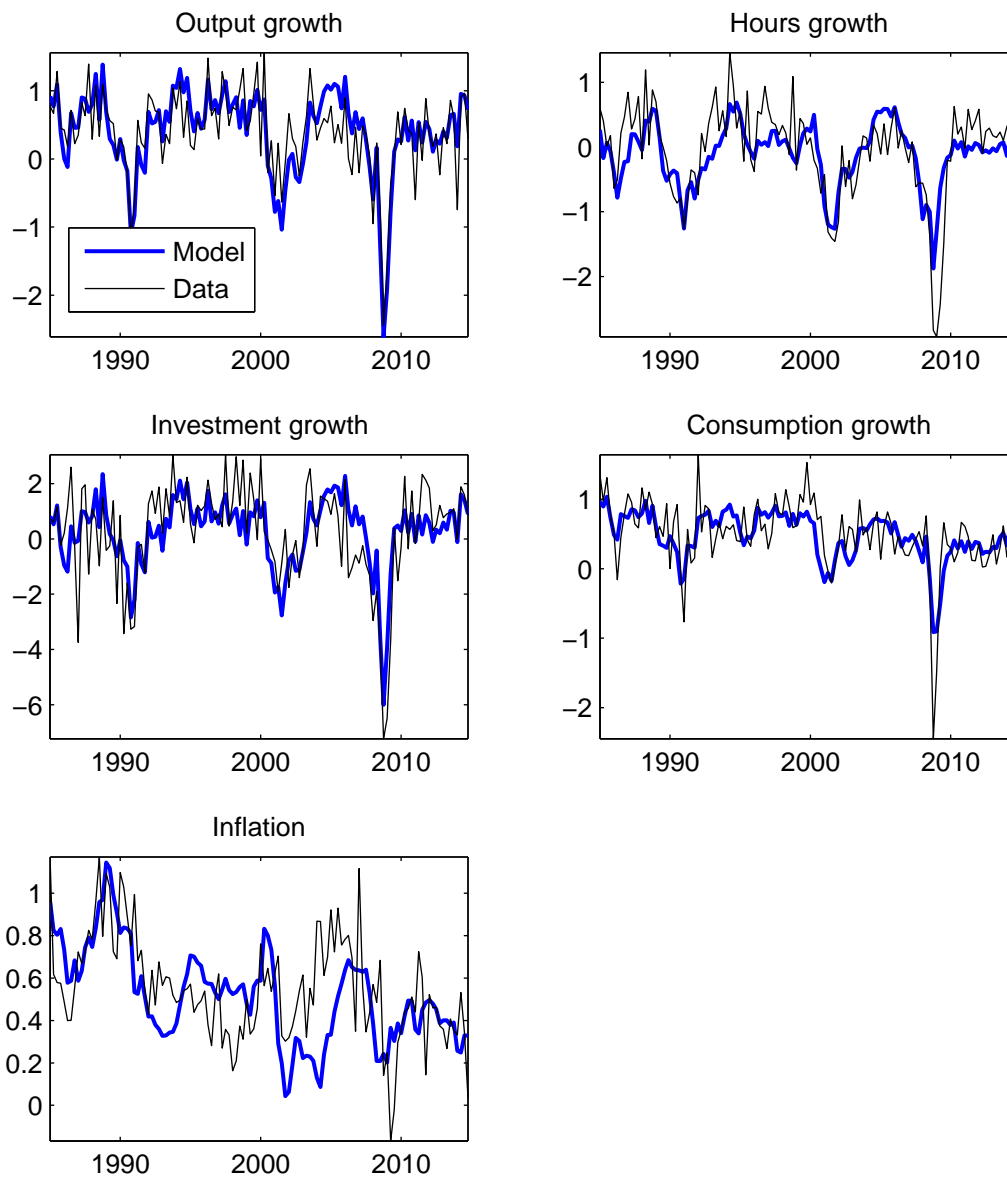
6.6 Additional figures from the estimated model

Figure 7: Data and predicted endogenous variables from the estimation: rational expectations model



Notes: Thick blue line ('Model') is the smoothed variable from the estimated model and the thin black line ('RE counterfactual') is the data. By construction, the difference between the model and the data is due to the measurement error.

Figure 8: Data and predicted endogenous variables from the estimation: model with ambiguity



Notes: Thick blue line ('Model') is the smoothed variable from the estimated model and the thin black line ('RE counterfactual') is the data. By construction, the difference between the model and the data is due to the measurement error.

Figure 9: Estimated impulse response to a TFP shock: rational expectations model




Figure 10: Estimated impulse response to a financial shock: rational expectations model

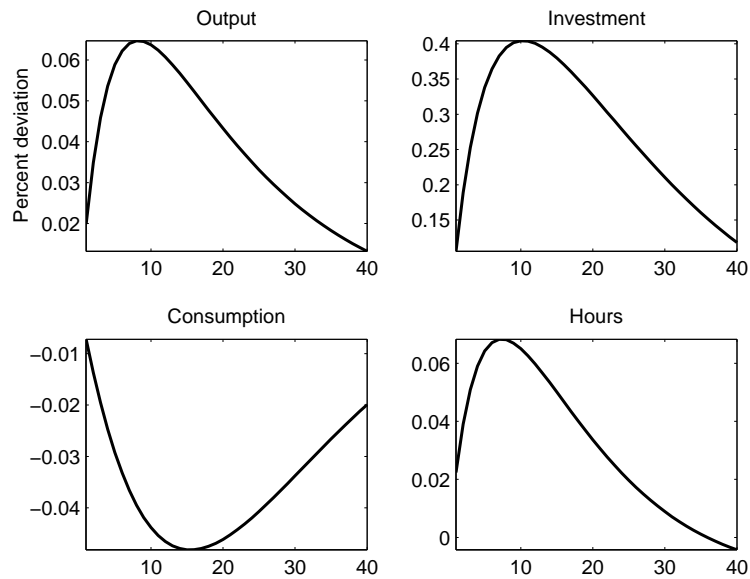
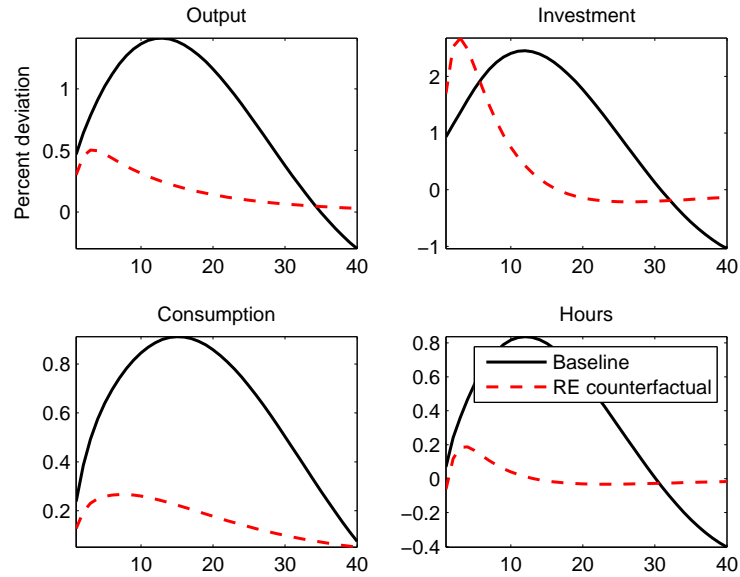


Figure 11: Estimated impulse response to a TFP shock: model with ambiguity



Notes: Black solid line ('Baseline') is the impulse response from our baseline model with endogenous uncertainty and the red dashed line ('RE counterfactual') is the response where we set the entropy constraint $\eta = 0$ while holding other parameters at their estimated values.